CiTiUS
Centro Singular de Investigación en Tecnoloxías da Información

arXiv:1609.05632v1 [cs.AI] 19 Sep 2016

# On the adoption of abductive reasoning for time series interpretation

T. Teijeiro and P. Félix

September 20, 2016

CAMPUS VIDA
USC UNIVERSIDADE DE SANTIAGO DE COMPOSTELA

**Abstract**

Time series interpretation aims to provide an explanation of what is observed in terms of its underlying processes. The present work is based on the assumption that common classification-based approaches to time series interpretation suffer from a set of inherent weaknesses whose ultimate cause lies in the monotonic nature of the deductive reasoning paradigm. In this document we propose a new approach to this problem based on the initial hypothesis that abductive reasoning properly accounts for the human ability to identify and characterize patterns appearing in a time series. The result of the interpretation is a set of conjectures in the form of observations, organized into an abstraction hierarchy, and explaining what has been observed. A knowledge-based framework and a set of algorithms for the interpretation task are provided, implementing a hypothesize-and-test cycle guided by an attentional mechanism. As a promising application domain, the interpretation of the electrocardiogram allows us to highlight the strengths of the present approach in comparison with traditional classification-based approaches.

# 1 Introduction

The interpretation and understanding of the behavior of a complex system involves the deployment of a cognitive apparatus aimed at guessing the processes and mechanisms underlying what is observed. The human ability to recognize patterns plays a paramount role as an instrument for pointing at the evidence that should require an explanation, by matching information from observations with information retrieved from memory. Classification naturally arises as a pattern recognition task defined as the assignment of observations to categories.

Let us first state precisely what is the problem under consideration: we wish to interpret the behavior of a complex system by measuring a physical quantity along time. This quantity is represented as a time series.

Artificial Intelligence community has devoted a lot of effort on different paradigms, strategies, methodologies and techniques for time series classification. Nonetheless, in spite of the wide range of proposals for building classifiers, either by eliciting domain knowledge or by induction from a set of observations, the resulting classifier behaves as a deductive system. The present work is premised on the assumption that some of the important weaknesses of this approach lie in its deductive nature, and that an abductive approach can address these shortcomings.

Let us remember that a deduction contains in its conclusions an information that is already implicitly contained in the premises, and thus is truth preserving. In this sense, a classifier ultimately assigns a label or a set of labels to each observation. This label can designate a process or a mechanism of the system being observed, but it is not more than a term that summarizes the premises satisfied by the observations. Conversely, abduction, or inference to the best explanation, is a form of inference that goes from data to a hypotheses that best explains or accounts for the data [15]. Abductive conclusions contain new information, not contained in the premises, capable of predicting new evidence, although they are fallible. Abductions are thus truth widening, and they can make the leap from the language of observations to the language of the underlying processes and mechanisms, responding to the aforementioned problem in a natural way [17]. For example, consider

a simple rule stating that if a patient experiences a sudden tachycardia and a decrease in blood pressure, then we can conclude that is suffering a shock due to the loss of blood volume. From a deductive perspective, *loss of blood volume* is just a name provided by the rule for the satisfaction of the two premises. However, from an abductive perspective, *loss of blood volume* is an explanatory hypothesis that expands the truth contained in the premises, enabling to predict additional consequences like for example skin paleness, faintness, dizziness or thirst.

A classifier is based on the assumption that the underlying processes or mechanisms are mutually exclusive. Superpositions and interactions of two or more processes, modifying their respective observation patterns, are excluded, and they must be represented by a new process, corresponding to a new category, different and usually unrelated to previous ones. Therefore, an artificial casuistry-based heuristics is adopted, increasing the complexity of the interpretation and reducing its adaptability to the variability of observations. In contrast, abduction infers a set of explanations as far as the available evidence does not allow us to identify the best one, and they are not incompatible with each other.

A better interpretation and understanding of a complex system should involve a larger number of observations, even future ones. Such interpretation is usually organized into a set of abstraction layers, where at each layer the temporal granularity of representation is reduced from below. A classification strategy provides an interpretation as the result of connecting a set of classifiers in cascade. Monotonicity of deduction entails a propagation of errors from the first abstraction layers upwards, narrowing the capability of making a proper interpretation as new abstraction layers are successively added. Following an abductive process instead, at each abstraction layer a best explanation hypothesis is conjectured from the data by the layer or layers below, within the context of information from above. Every observation made by abstracting from raw data upwards is a conjecture, and non monotonicity of abduction supports retracting any observation at any abstraction layer in search for the best global explanation. Thus, bottom-up and top-down processing complement one another and jointly provide a result.

In a classifier the truth of the conclusion follows from the truth of all the premises, and missing data usually demand for some imputation strategy that results in a conjecture: a sort of abducing to go on deducing. In contrast, an abductive interpretation is posed as a hypothesize-and-test cycle, where the evidence is incrementally added to reasoning as soon as it is available. In this sense, an abductive interpretation adapts to time-varying properties of the complex system, providing the current result of the interpretation process at any given time.

Abduction has been primarily proposed for diagnostic tasks [24, 6], but also for question answering [11], language understanding [16], story comprehension [3], image understanding [27] or plan recognition [20], amongst others. Some researches have proposed that perception could rely on some form of abduction. Even though abductive reasoning has been proved to be NP-complete, a compiled form of abduction, based on a set of prestored hypotheses could narrow the generation of hypotheses [17]. The present work takes this assumption as a starting point and proposes a model-based abductive framework for time series interpretation supported on a set of temporal abstraction patterns. An abstraction pattern represents a set of constraints that must be satisfied by some evidence for being interpreted as the hypothetical observation of a certain process, together with an observation procedure providing a set of measurements for the features of the conjectured observation.

A set of algorithms is devised in order to achieve the best explanation through a process of successive abstraction from raw data, by means of a hypothesize-and-test strategy.

The present work includes some examples from electrocardiography, the recording on the body surface of the electrical activity of the heart as it changes with time. The electrocardiogram (ECG) is the primary method for the study and diagnosis of the cardiac disease, since the processes involved in cardiac physiology manifest in characteristic temporal patterns on the ECG tracing. In other words, a correct reading of the ECG has the potential to provide valuable insight into cardiac phenomena. Learning to interpret the ECG involves the acquisition of perceptual skills from an extensive bibliography with interpretation criteria and worked examples. In particular, pattern recognition is specially important in order to build a bottom-up representation of the cardiac phenomena in multiple abstraction levels. This has encouraged extensive research on classification techniques for interpreting the ECG but in spite of all the efforts, it is considered an open problem. We shall try to demonstrate that the problem is in the nature of deduction itself.

The rest of this paper is outlined as follows: Section 2 introduces the main concepts and terminology used in the paper in an informal and intuitive way. Afterwards, in sections 3, 4 and 5 we formally describe all the components of the interpretation framework, including the knowledge representation model and the algorithms to obtain effective interpretations in an affordable time. Section 6 illustrates the capabilities of the framework to overcome some the most important shortcomings of deductive classifiers. Finally, in section 7 we discuss the properties of the model compared with other related approaches, and we draw some conclusions.

# 2 Interpretation as a process-guessing task

We propose a knowledge-based interpretation framework upon the principles of abductive reasoning, on the basis of a strategy of hypothesis formation and testing. Taking as starting point a time series of physical measurements, a set of observations are guessed as conjectures of the underlying processes, through successive levels of abstraction. Each new observation will be generated from previous levels, as underlying processes aggregate, superimpose or concatenate to form more complex processes with greater duration and scope, being organized in an abstraction hierarchy.

The knowledge of the domain is described as a set of abstraction patterns as follows:

$$h_\psi((\mathbf{A}_h, T_h) = \pi(\mathbf{A}_1, T_1, ..., \mathbf{A}_n, T_n)) \;\; abstracts \;\; m_1(\mathbf{A}_1, T_1), ..., m_n(\mathbf{A}_n, T_n) \;\; \{C(\mathbf{A}_h, T_h, \mathbf{A}_1, T_1, ..., \mathbf{A}_n, T_n)\}$$

where $h_\psi(\mathbf{A}_h, T_h)$ is an observable of the domain playing the role of a hypothesis on the observation of an underlying process $\psi$, where $\mathbf{A}_h$ represents a tuple of attributes, and its temporal support is represented as an instant $T_h$ for the sake of simplicity, but it can also be an interval; $m_1(\mathbf{A}_1, T_1)$,...,$m_n(\mathbf{A}_n, T_n)$ is a set of observables of the domain playing the role of the evidence that suggests the observation of $h_\psi$; $C$ is a set of constraints among the variables involved in the abstraction pattern; constraints in $C$ are interpreted as necessary conditions in order for the evidence $m_1(\mathbf{A}_1, T_1)$,...,$m_n(\mathbf{A}_n, T_n)$ to be abstracted into $h_\psi(\mathbf{A}_h, T_h)$; $\pi(\mathbf{A}_1, T_1, ..., \mathbf{A}_n, T_n)$ is an abstraction procedure that gives as a result an observation of $h_\psi(\mathbf{A}_h, T_h)$ from a set of observations for the set $m_1(\mathbf{A}_1, T_1)$,...,$m_n(\mathbf{A}_n, T_n)$.

To illustrate this concept, consider the sequence of observations in figure 1. Each one of these observations is an instance of an observable we call *point* ($p$), that is represented as $p(\mathbf{A} = \langle V \rangle, T)$, where $T$ determines the temporal location of the observation and $V$ is a value attribute.

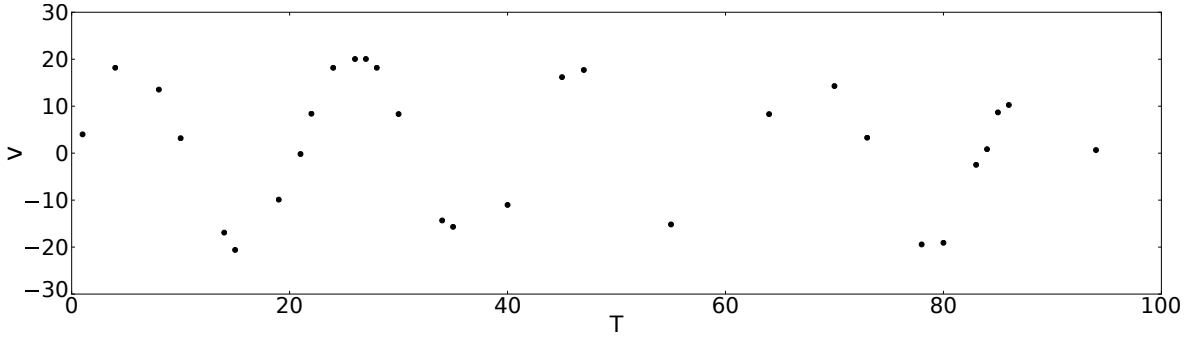


Figure 1: Initial temporal observations

If we visually analyze these observations, we may hypothesize the presence of an underlying sinusoidal process. Let us define an observable $sinus$ for such a sinusoidal process, with two attributes: the amplitude of the process ($\alpha$) and its frequency ($\omega$). The necessary knowledge to conjecture this hypothesis is collected in the following abstraction pattern:

$$h_{sinus}(\mathbf{A}_h = \langle \alpha, \omega \rangle = \pi(V_1, T_1, \ldots, V_n, T_n)) \;\; abstracts \;\; p(V_1, T_1), ..., p(V_n, T_n) \\ \{C(V_1, T_1, ..., V_n, T_n)\}$$

We can estimate the characteristics $\langle \alpha, \omega \rangle$ of this process by a simple abstraction procedure $\pi$, such that $\alpha = max(|V_i|)$, for $1 \leq i \leq n$, that is, the amplitude $\alpha$ is obtained as the maximum absolute value of the observations; and $\omega = \pi / mean(T_j^{peak} - T_{j-1}^{peak})$, where $(V_j^{peak} = V_k, T_j^{peak} = T_k) \wedge sign(V_k - V_{k-1}) \neq sign(V_{k+1} - V_k)$, that is, the frequency $\omega$ is obtained as the inverse of the mean temporal separation between consecutive peaks in the sequence of observations.

We can impose the following constraint $C(V_1, T_1, ..., V_n, T_n)$ for every pair $(V_i, T_i)$ in the sequence:

$$max(|\alpha \cdot sin(\omega \cdot T_i) - V_i|) \leq \epsilon,$$

This constraint provides a model of a sinusoidal process and a measure of how well it fits a set of observations, by means of a maximum error $\epsilon$. Figure 2 shows the continuous representation of the abstracted process and its relation with the original observations. A value of $\alpha/3$ has been chosen for $\epsilon$.

Of course, different abstraction methods can be devised in order to estimate the same or different characteristics of the process being guessed. These methods can provide one or multiple valid estimations in terms of their consistency with the above mentioned necessary constraints. And different processes can be guessed from the same set of observations, all of them being valid in terms of their consistency. Hence, we can need further criteria in order to rank the set of interpretations.

This simple example summarizes the common approach to interpretation of experimental results in science and technology, when the knowledge is available as a model or a set of

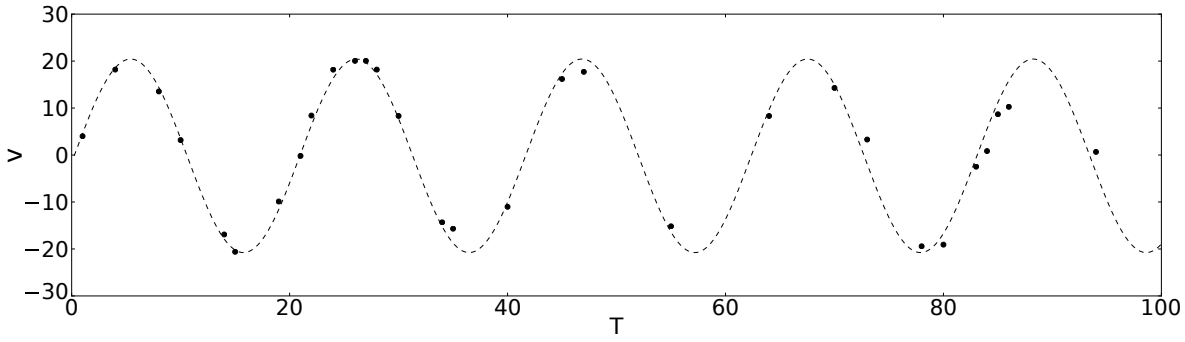


Figure 2: Abstracted sinusoidal process

models. The challenge is to assume that knowledge is not available in an analytical but in a declarative form, as a pattern or a set of patterns, and the interpretation task is expected to mimic some mechanisms of human perception.

# 3 Definitions

In this section we formally define the main pieces of our interpretation framework: *observables* and *observations* for representing the behavior of the system under study, and *abstraction patterns* for representing the knowledge about this system.

## 3.1 Representation entities

An *observation* is the result of applying an observation procedure. Such procedure makes it possible to observe something with the quality of being *observable*. We call $\mathcal{Q} = \{q_0, q_1, ..., q_n\}$ to the set of observables of a particular domain.

**Definition 1.** We define an **observable** as a tuple $q = \langle \psi, \mathbf{A}, T^b, T^e \rangle$, where $\psi$ is a name representing the underlying process being observable, $\mathbf{A} = (A_1, ..., A_{n_q})$ is a set of attributes to be valued, and $T^b$ and $T^e$ are two temporal variables representing the beginning and the end of the observable.

We call $V(A_i)$ to the domain of possible values for the attribute $A_i$. For any observable, we implicitly assume the constraint $T^b < T^e$. In case of an instantaneous observable, then it is represented as $q = \langle \psi, \mathbf{A}, T \rangle$. Some observables can be dually represented from the temporal perspective, as an observable supported by a temporal interval or as an observable supported by a temporal instant, according to the task to be carried out. A paradigmatic example is found in representing the heart beat, since it can be represented as a domain entity with a temporal extension comprising its constituent waves, and it can also be represented as an instantaneous entity for measuring the heart rate.

**Example 3.1.** *In the ECG signal several distinctive waveforms can be identified, corresponding to the electrical activation-recovery cycle of the different heart chambers. The so-called P wave represents the activation of the atria, and is the first wave of the cardiac cycle. The next group of waves recorded is the QRS complex, representing the simultaneous activation of the right and left ventricles. Finally, the wave that represents the ventricular*

*recovery is called the T wave. Together, these waveforms devise the characteristic pattern of the heart cycle, which is repeated in a normal situation with every beat [31]. An example of a common ECG strip is shown in figure 3.*

*According to this description, the observable* $q_{Pw} = \langle \texttt{atrial_activation}, \texttt{amplitude}, T^b, T^e \rangle$ *represents a P wave resulting from an atrial activation process with an unknown amplitude localized in a still unknown temporal interval.*

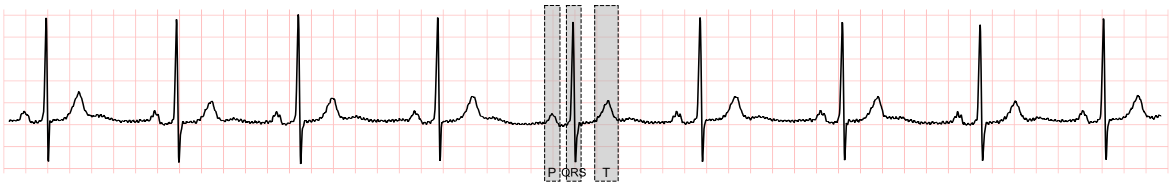


Figure 3: Example of the ECG basic waveforms. [Source: MIT-BIH arrhythmia DB [14], recording: 123, between 12:11.900 and 12:22.400]

**Definition 2.** We define an **observation** as a tuple $o = \langle q, \mathbf{v}, t^b, t^e \rangle$, the result of applying an observation procedure over an observable $q$, where $\mathbf{v} = (v_1, ..., v_{n_q})$ is a set of attribute values, such that $A_1 = v_1$,...,$A_{n_q} = v_{n_q}$; and $t^b$ and $t^e$ are two precise instants, such that $T^b = t^b$ and $T^e = t^e$, limiting the beginning and the end of the observation.

**Example 3.2.** *The tuple* $o = \langle q_{Pw}, \texttt{amplitude} = 0.17\texttt{mV}, \textit{12:16.977}, \textit{12:17.094} \rangle$ *represents the particular occurrence of the P wave observable highlighted in figure 3.*

We call $\mathcal{O}$ to the set of observations available for the observables in $\mathcal{Q}$. Now we define some relations involving observables and observations, that will be useful to describe some properties and constraints of the domain concepts, as well as to temporally arrange the interpretation process.

**Definition 3.** Given a set of observables $\mathcal{Q}$, a **generalization relation** can be defined between two different observables $q = \langle \psi, \mathbf{A}, T^b, T^e \rangle$ and $q' = \langle \psi', \mathbf{A}', T'^b, T'^e \rangle$, denoted by $q'$ `is a` $q$, meaning that $q$ generalizes $q'$. Every generalization satisfies that $V(A'_i) \subseteq V(A_i)\ \forall i \in [1, n_q]$, and that $\mathbf{A} \subseteq \mathbf{A}'$.

The generalization relation is reflexive, antisymmetric and transitive. The inverse of a generalization relation is a specification relation. From a logical perspective, a generalization relation can be read as an implication $q' \rightarrow q$, meaning that $q'$ is more specific than $q$. It holds that every observation $o = \langle q', \mathbf{v}, t^b, t^e \rangle$ of the observable $q'$ is also an observation of $q$.

**Example 3.3.** *A common example of a generalization relation can be defined from a domain partition of an attribute. For example, the observable* $q_1 = \langle \texttt{Sinus_Rhythm}, (\texttt{RR} \in [200ms, 4000ms]), T^b, T^e \rangle$ *is a generalization of the observables* $q_2 = \langle \texttt{Sinus_Tachycardia}, (\texttt{RR} \in [200ms, 600ms]), T^b, T^e \rangle$, $q_3 = \langle \texttt{Normal_Rhythm}, (\texttt{RR} \in [600ms, 1000ms]), T^b, T^e \rangle$ *and* $q_4 = \langle \texttt{Sinus_Bradycardia}, (\texttt{RR} \in [1000ms, 4000ms]), T^b, T^e \rangle$. *The* `RR` *attribute represents the measure, in milliseconds, of the mean distance between consecutive beats, while* $q_2, q_3$ *and* $q_4$ *represent the normal cardiac rhythm denominations according to the heart rate [31].*

**Definition 4.** Given a set of observables $\mathcal{Q}$, an **exclusion relation** can be defined between two different observables $q = \langle \psi, \mathbf{A}, T^b, T^e \rangle$ and $q' = \langle \psi', \mathbf{A}', T'^b, T'^e \rangle$, denoted by $q$ `excludes` $q'$, meaning that two different observables are mutually exclusive if and only if they cannot be observed over two properly overlapping intervals, that is, $T^e < T'^b \vee T'^e < T^b$.

The exclusion relation is reflexive, symmetric and transitive. Hence, every observation $o_i = \langle q_i, \mathbf{v}_i, t_i^b, t_i^e \rangle$ excludes a different observation $o_j = \langle q_j, \mathbf{v}_j, t_j^b, t_j^e \rangle$ of the same observable at the same time, that is, no attribute may take different values simultaneously: $\forall o_i, o_j \in \mathcal{O}, (q_i = q_j) \wedge (\mathbf{v}_i \neq \mathbf{v}_j) \Rightarrow (t_i^e < t_j^b) \vee (t_j^e < t_i^b)$. Furthermore, the observation of the same values is also excluded over two properly overlapping intervals, assuming that the evidence explained by the later one will be subsumed in the first one, which will be temporally extended. A subsumption procedure will be described later in this paper. The generalization relation may hidden the exclusion relation, since the observation of multiple non exclusive observables over properly overlapping intervals can appear as the multiple observation of a single observable which is a generalization of the above: $\forall o_i = \langle q, \mathbf{v}, t_i^b, t_i^e \rangle, o_j = \langle q, \mathbf{v}, t_j^b, t_j^e \rangle \in \mathcal{O} \mid (t_i^e > t_j^b) \wedge (t_i^e < t_j^e) \Rightarrow \exists\ q_i', q_j' \in \mathcal{Q} \mid o_i = \langle q_i', \mathbf{v}, t_i^b, t_i^e \rangle \wedge o_j = \langle q_j', \mathbf{v}, t_j^b, t_j^e \rangle \wedge q_i'$ `is a` $q \wedge q_j'$ `is a` $q \wedge q_i' \neg$`excludes` $q_j'$.

**Example 3.4.** *The R-on-T phenomenon is the superimposition of an ectopic beat on the T wave of a preceding beat [31]. There is an overlapping of two observations of the observable wave as a generalization of the R and T waves.* An example is shown in figure 4.

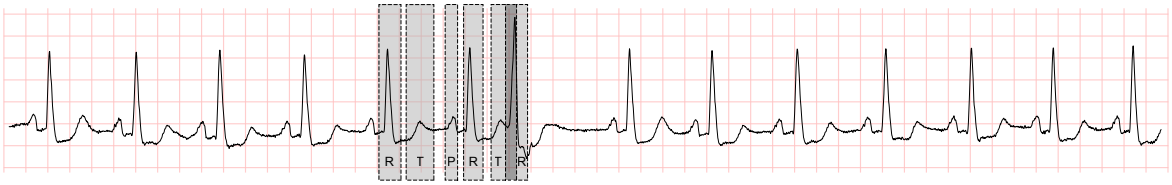


Figure 4: Example of the R-on-T phenomenon. [Source: MIT-BIH arrhythmia DB, recording: 105, between 00:10.800 and 00:21.300]

We call $\mathcal{O}$ to the set of observations requiring interpretation. In order to index the set of observations they will be represented as a sequence by defining an order relation between them. This ordering aims to prioritize the interpretation of the observations as they appear.

**Definition 5.** Let $<$ be an **order relation** between two observations $o_i = \langle q_i, \mathbf{v}_i, t_i^b, t_i^e \rangle$ and $o_j = \langle q_j, \mathbf{v}_j, t_j^b, t_j^e \rangle$ such that $(o_i < o_j) \Leftrightarrow (t_i^b < t_j^b) \vee ((t_i^b = t_j^b) \wedge (t_i^e < t_j^e)) \vee ((t_i^b = t_j^b) \wedge (t_i^e = t_j^e) \wedge (q_i < q_j))$, assuming a lexicographical order between observable names.

A sequence of observations is an ordered set of observations $\mathcal{O} = (o_1, ..., o_i, ...)$ where for all $i < j$ then $o_i < o_j$. Every subset of a sequence of observations is also a sequence. The temporal arrangement of observations in a sequence is restricted when the observations correspond to the same observable:

**Definition 6.** We define a **q-sequence of observations** as a sequence of observations for the observable $q$, denoted as $O(q)$, where any pair of observations $o_i = \langle q, \mathbf{v}_i, t_i^b, t_i^e \rangle$ and $o_j = \langle q, \mathbf{v}_j, t_j^b, t_j^e \rangle$ holds that $t_i^e < t_j^b$ or $t_j^e < t_i^b$.

By $\texttt{succ}(o_i)$ we denote the successor of the observation $o_i$ in the sequence $\mathcal{O}$, according to the precedence ordering $<$. By q-$\texttt{succ}(o_i)$ we denote the successor of the observation $o_i \in O(q)$ in its q-sequence $O(q)$. Conversely to this notation, we denote by $q(o_i)$ the observable corresponding to the $o_i$ observation.

## 3.2 Abstraction Patterns

Abstraction patterns are the knowledge representation primitives of this interpretation framework. An abstraction pattern allows to conjecture the occurrence of an observable from some other observables that appear as findings occurring in a particular temporal arrangement. Formally, an abstraction pattern is defined as follows:

**Definition 7.** An **abstraction pattern** $P = \langle q_h, M_P, C_P, \pi_P \rangle$ consists of a hypothesis observable $q_h$, a set of findings $M_P = \{m_1, \ldots, m_n\}$, a set of constraints $C_P(\mathbf{A}_h, T_h, \mathbf{A}_1, T_1, ..., \mathbf{A}_n, T_n)$ checked according to an ordering compatible with their dependencies, and an observation procedure $\pi_P : (\mathbf{A}_1, T_1, ..., \mathbf{A}_n, T_n) \rightarrow O(q_h)$.

An abstraction pattern establishes the set of conditions and calculations for producing a new observation $o_h \in O(q_h)$ as a conjecture from the observation of a set of findings. Each finding numbers the multiple occurrences of a given observable in the pattern. We call $M_P^q = \{m_1^q, m_2^q, ..., m_i^q = \langle \psi, \mathbf{A}_i, T_i^b, T_i^e \rangle, ...\}$ to the set of findings of the observable $q$ in $P$. The interpretation procedure will choose, as we will see later, among the available observations for the observable $q$ satisfying the constraints $C_P$, those to be assigned to the findings in $M_P^q$ in order to calculate $o_h$.

The set of findings $M_P$ is divided into two sets $A_P$ and $E_P$, being $A_P \cap E_P = \varnothing$, where $A_P$ is the set of findings that is said to be abstracted in $o_h$, and $E_P$ is the set of findings that constitute the observation environment of $o_h$, that is, the set of findings needed to properly conjecture $o_h$, but which are not synthesized in $o_h$.

A temporal covering assumption can be made as a default assumption on a hypothesis with respect to those findings appearing in an abstraction pattern:

**Default Assumption 1.** (*Temporal covering) Given an abstraction pattern $P$, it holds that $T_h^b \leq T_m^b$ and $T_m^e \leq T_h^e$, for all $m \in A_P \subseteq M_P$.*

The temporal covering assumption will allow us to define the exclusiveness of an interpretation, as the impossibility of including competing abstractions in the same interpretation.

**Example 3.5.** *In electrocardiography, a wave is a discernible deviation from a horizontal reference line called baseline, where at least two opposite slopes can be identified. The term discernible means that both the amplitude and the duration of the deviation must exceed some minimum values, that are agreed to be 20 μV and 6 ms, respectively [7]. A wave can be completely described by its amplitude ($A$), duration ($T^e - T^b$), voltage polarity ($V \in \{+, -\}$) and its turning point $T^{tp} \in [T^b, T^e]$, resulting in the following observable:*

$$q_{wave} = \langle \texttt{electrical_activity}, (A, V, T^{tp}), T^b, T^e \rangle$$

*Consider the following abstraction pattern:*

$$P_{wave} = \langle q_{wave}, M_P = \{m_0^{ECG}, \ldots, m_n^{ECG}\}, C_{P_{wave}}, \texttt{wave_observation}() \rangle$$

*where a set of temporal constraints can be established between the temporal variables,* $C_1 \equiv \langle T^e - T^b \geq 6ms \rangle$, $C_2 \equiv \langle T^b = T_{m_1} \rangle$ *and* $C_3 \equiv \langle T^e = T_{m_{n-1}} \rangle$ *, and another set of constraints limit the amplitude and slope changes of the samples included in a wave,* $C_4 \equiv \langle sign(m_1 - m_0) \neq sign(m_2 - m_1) \rangle$, $C_5 \equiv \langle sign(m_n - m_{n-1}) \neq sign(m_{n-1} - m_{n-2}) \rangle$, $C_6 \equiv \langle sign(m_{tp} - m_{tp-1}) = -sign(m_{tp+1} - m_{tp}) \rangle$ *and* $C_7 \equiv \langle min\{|m_{tp} - m_1|, |m_{tp} - m_{n-1}|\} \geq 20\mu V \rangle$.

*Once a set of ECG samples has satisfied these constraints, they support the observation of a wave, and the observation procedure provides a value for its parameters, for example,* $V = sign(m_{tp} - m_1)$, $A = max\{|m_{tp} - m_1|, |m_{tp} - m_{n-1}|\}$, *and* $T^{tp} = arg\,min_k\{m_k | 1 \leq k \leq n-1\}$, *if* $m_1 < m_0$, *or* $T^{tp} = arg\,max_k\{m_k | 1 \leq k \leq n-1\}$, *if* $m_1 > m_0$.

*The onset and end of a wave are set to the time of the second and second last samples, since* $m_0$ *and* $m_n$ *are environment observations used to check the presence of a slope change just before and after the wave, so* $E_{P_{wave}} = \{m_0, m_n\}$, *and* $A_{P_{wave}} = \{m_1, \ldots, m_{n-1}\}$.

According to the definition, an abstraction pattern is defined over a fixed set of evidence findings $M_P$. In general, however, an abstraction involves an undetermined number of pieces of evidence -in the case of an ECG wave the number of samples-. Hence we provide a procedure for dynamically generating abstraction patterns, based on the formal language theory. The set $\mathcal{Q}$ of observables can be considered as an alphabet. Given an alphabet $\mathcal{Q}$, a couple of special symbols $\varnothing$ (empty set), and $\lambda$ (empty string), and operators | (union), $\cdot$ (concatenation), and $*$ (kleene closure), a formal grammar $G$ denotes a pattern of symbols of the alphabet, describing a language $L(G) \subseteq \mathcal{Q}^*$, as a subset of the set of possible strings of symbols of the alphabet.

Let $G^{ap}$ be the class of formal attributed grammars of abstraction patterns. An *abstraction grammar* $G \in G^{ap}$ is syntactically defined as a tuple $(V_N, V_T, H, R)$. We adopt a simple expressiveness for the production rules in $R$:

$$\begin{aligned} H = q_h &\rightarrow q[c]D \\ D &\rightarrow q[c]F \mid q[c] \mid \lambda \end{aligned}$$

$H = q_h$ is the initial symbol of the grammar, and it plays the role of the hypothesis guessed by the patterns generated by $G$. A single grammar allows us to describe all the different ways of abstracting an observable $q_h$. $V_N$ is the set of non-terminal symbols of the grammar, satisfying $H \notin V_N$ since an observable cannot explain itself. $V_T$ is the set of terminal symbols of the grammar, gathering together: a set of observables $Q_G \subseteq \mathcal{Q}$ that can be abstracted by the hypothesis, and a set of constraint descriptions $[c]$ between all the observables.

Given a grammar $G \in G^{ap}$, we provide a constructive method for representing a set of abstraction patterns $P_G = \{P_1, ..., P_i, ...\}$. An abstraction pattern aims to describe a set of observables as findings to be abstracted. When these findings appear, possibly repeatedly, under certain constraints they can be abstracted in a new observable $q_h$. $P_G$ gathers the set of abstraction patterns that share the same observable $q_h$ as hypothesis, so they represent the different ways to conjecture the presence of $q_h$.

Besides describing all the possible sets of findings that may be abstracted by $q_h$, grammars in $G^{ap}$ must also define the set of constraints $C$, the observation procedure $\pi$, and the membership of every finding to $A_P$ or $E_P$. For this, every grammar $G \in G^{ap}$ is semantically extended with a set of attributes, so $G = ((V_N, V_T, H, R), B, BR)$, where $B$ associates

each grammar symbol $\alpha \in V_N \cup V_T$ with a set of attributes, and $BR$ associates each rule $r \in R$ with a set of attribute computation rules. The set of attribute computation rules comprises two types of rules: 1) the semantic expression of the constraint descriptions $[c]$ between all the observables, leading to the set of constraints $C$ that must be satisfied by those findings involved in an abstraction pattern; and 2) the observation procedure $\pi$ tied to the symbol $q_h$, that establishes a calculation for producing an observation $o_h \in O(q_h)$ from observations of the findings $M_P$ involved in the abstraction pattern.

In the following, we devise a constructive method to build abstraction patterns from a grammar description $G \in G^{ap}$. Using this method, the application of every production incrementally adds a new observable as a finding and a set of constraints among this finding and those generated previously:

1. The start symbol $H$ entails initializing an abstraction pattern with the hypothesis:

$$P \leftarrow \langle q_h, M_P = \varnothing, C_P = C(T_h^b, T_h^e, \mathbf{A}_h), \pi_P(\mathbf{A}_h)\rangle$$

2. All those productions of the form $H = q_h \rightarrow q[c]D$ entail:

$$P \leftarrow \langle q_h, M_P = \{m_1^q\}, C_P \cup C(T_1^b, T_1^e, \mathbf{A}_1), \pi_P(\mathbf{A}_h, \mathbf{A}_1)\rangle$$

3. All those productions of the form $D \rightarrow q[c]F \mid q[c]$ entail:

$$P \leftarrow \langle q_h, M_P \cup \{m_k^q\}, C_P \cup C(T_k^b, T_k^e, \mathbf{A}_k), \pi_P(\mathbf{A}_h, \mathbf{A}_1, \ldots, \mathbf{A}_k)\rangle$$

4. All those productions of the form $D \rightarrow \lambda$ entail:

$$P \leftarrow P$$

This constructive method enables the incremental addition of new constraints as new findings are included in the representation of the abstraction pattern, providing a dynamic mechanism of knowledge assembling by language generation. Moreover, it is possible to design an adaptive observation procedure as new evidence becomes available, since $\pi_P(\mathbf{A}_h, \mathbf{A}_1, \ldots, \mathbf{A}_n)$ can be different at each step.

In case the constraints description $[c]$ is omitted in a production, it is assumed the 'after' temporal relationship between the new finding and the set of previous findings. For instance, all those productions of the form $C \rightarrow qF \mid q$ entail:

$$P \leftarrow \langle q_H, M_P \cup \{m_k^q\}, C_P \cup \{T_i \leq T_k^b \mid m_i \in M_P\}, \pi_P(\mathbf{A}_h, ..., \mathbf{A}_k)\rangle$$

Hence, in the absence of any temporal description, an increasing temporal order among consecutive findings in every abstraction pattern is assumed. Moreover, every temporal description must be consistent with this temporal order.

According to the limitation imposed on observations of the same observable, preventing two different observations from occurring at the same time, an additional constraint is added on any two findings of the same observable, and thus $\forall m_i^q, m_j^q \in M_P, (T_i^e < T_j^b \vee T_j^e < T_i^b)$.

Below we show some examples of abstraction pattern grammars modeling common knowledge in electrocardiography, in order to illustrate the expressiveness of the $G^{ap}$ grammars.

**Example 3.6.** *The grammar $G_N = (V_N, V_T, H = q_N, R)$ is designed to generate an abstraction pattern for a normal cardiac cycle, including the descriptions of common durations and intervals [31]. In this grammar, $V_N = \{D, E\}$, $V_T = \{q_{Pw}, q_{QRS}, q_{Tw}\}$, and $R$ is given by:*

$$\begin{aligned} H = q_N &\rightarrow q_{Pw}[c_1]D \\ D &\rightarrow q_{QRS}[c_2]E \\ E &\rightarrow q_{Tw}[c_3] \end{aligned}$$

*where a set of temporal constraints can be established among the constituent waves: $c_{11} \equiv \langle T^b_N = T^b_{Pw} \rangle$ and $c_{12} \equiv \langle 50ms \leq T^e_{Pw} - T^b_{Pw} \leq 120ms \rangle$ for the first production rule; $c_{21} \equiv \langle 50ms \leq T^e_{QRS} - T^b_{QRS} \leq 150ms \rangle$ and $c_{22} \equiv \langle 100ms \leq T^b_{QRS} - T^b_{Pw} \leq 210ms \rangle$ for the second production rule; and $c_{31} \equiv \langle 80ms \leq T^b_{Tw} - T^e_{QRS} \leq 120ms \rangle$, $c_{32} \equiv \langle T^e_{Tw} - T^b_{QRS} \leq 520ms \rangle$ and $c_{33} \equiv \langle T^e_N = T^e_{Tw} \rangle$ for the last production rule.*

*Along the execution of these production rules an abstraction pattern is generated:*

1. *The start symbol $H$ entails initializing the observable pattern:*

$$P \leftarrow \langle q_N, M_P = \varnothing, C_P = \varnothing, \pi_P = \varnothing \rangle$$

2. *The production $H = q_N \rightarrow q_{Pw}[c_1]D$ entails:*

$$P \leftarrow \langle q_N, M_P = \{m_1 = q_{Pw}\}, C_P = \{c_{11}, c_{12}\}, \pi_P \rangle$$

3. *The production $D \rightarrow q_{QRS}[c_2]E$ entails:*

$$P \leftarrow \langle q_N, M_P \cup \{m_2 = q_{QRS}\}, C_P \cup \{c_{21}, c_{22}\}, \pi_P \rangle$$

4. *The production $E \rightarrow q_{Tw}[c_3]$ entails:*

$$P \leftarrow \langle q_N, M_P \cup \{m_3 = q_{Tw}\}, C_P \cup \{c_{31}, c_{32}, c_{33}\}, \pi_P \rangle$$

*This grammar generates a single abstraction pattern, which allows us to interpret a sequence of a P wave, a QRS complex, and a T wave as a normal cardiac cycle, where some additional temporal constraints are required. The abstraction pattern allows us to specify the temporal location of the hypothesis, a normal cardiac cycle, from the location of the findings. In this case, an observation procedure $\pi_P$ is not necessary since the attributes of the hypothesis are completely determined by the constraints in the grammar, not requiring any additional calculus.*

The next example shows the ability of an abstraction grammar to dynamically generate abstraction patterns with an undefined number of findings.

**Example 3.7.** *A bigeminy is a heart arrhythmia in which there is a continuous alternation of long and short heart beats. Most often this is due to ectopic heart beats occurring so frequently that there is one after each normal beat, typically premature ventricular contractions (PVCs) [31]. For example, a normal beat is shortly followed by a PVC, then followed by a pause. The normal beat then returns, only to be followed by another PVC.*

*The grammar $G_{VB} = (V_N, V_T, H = q_{VB}, R)$ is designed to generate a set of abstraction patterns for ventricular bigeminy, where $V_N = \{D, E, F\}$, $V_T = \{q_N, q_V\}$, and $R$ is given by:*

$$\begin{aligned} H = q_{VB} &\rightarrow q_N[c_1]D \\ D &\rightarrow q_V[c_2]E \\ E &\rightarrow q_N[c_3]F \\ F &\rightarrow q_V[c_4]E \mid q_V[c_4, c_5] \end{aligned}$$

*The set of constraints is defined as following: $c_1 \equiv \langle T^b_{VB} = T_1 \rangle$, $c_2 \equiv \langle 200ms \leq T_2 - T_1 \leq 800ms \rangle$, $c_3 \equiv \langle 1.5 \cdot 200ms \leq T_k - T_{k-1} \leq 4 \cdot 800ms \rangle$, $c_4 \equiv \langle 200ms \leq T_{k+1} - T_k \leq 800ms \rangle$ and $c_5 \equiv \langle T^e_{VB} = T_n \rangle$.*

*For simplicity, we have referenced each $N$ and $V$ heart beats with a single temporal variable, so $T_i$ represents the time point of the i-th heart beat, being a normal beat if $i$ is odd, and a PVC if $i$ is even. Along the execution of these production rules an unbounded sequence of alternating normal and premature ventricular QRS complexes is generated, what is just described as ventricular bigeminy. Note that in terms of the $\{N, V\}$ symbols the $G_{VB}$ grammar is equivalent to the regular expression $NV(NV)+$. The iterative building of abstraction patterns according to this grammar is as follows:*

1. *The start symbol $H$ entails initializing the pattern:*

$$P \leftarrow \langle q_{VB}, M_P = \varnothing, C_P = \varnothing, \pi_P = \varnothing \rangle$$

2. *The production $H = q_{VB} \rightarrow q_N[c_1]D$ entails:*

$$P \leftarrow \langle q_{VB}, M_P = \{m_1 = q_N\}, C_P \cup \{c_1\}, \pi_P \rangle$$

3. *The production $D \rightarrow q_V[c_2]E$ entails:*

$$P \leftarrow \langle q_{VB}, M_P \cup \{m_2 = q_V\}, C_P \cup \{c_2\}, \pi_P \rangle$$

4. *The production $E \rightarrow q_N[c_3]F$ entails, in the $k$ repetition:*

$$P \leftarrow \langle q_{VB}, M_P \cup \{m_k = q_N\}, C_P \cup \{c_3\}, \pi_P \rangle$$

5. *The production $F \rightarrow q_V[c_4]E$ entails, in the $k$ repetition:*

$$P \leftarrow \langle q_{VB}, M_P \cup \{m_{k+1} = q_V\}, C_P \cup \{c_4\}, \pi_P \rangle$$

6. *The production $F \rightarrow q_V[c_4, c_5]$ entails finishing the pattern with $n$ observations:*

$$P \leftarrow \langle q_{VB}, M_P \cup \{m_n = q_V\}, C_P \cup \{c_4, c_5\}, \pi_P \rangle$$

*In this example, as in 3.6, an observation procedure $\pi_P$ is not necessary, since the constraints in the grammar completely determine the temporal endpoints of the hypothesis and there are no more attributes to be valued. Figure 5 shows an example of a ventricular bigeminy pattern.*

Figure 5: Ventricular bigeminy example. [Source: MIT-BIH arrhythmia DB, recording: 106, between 25:06.350 and 25:16.850]

# 4 An interpretation framework

In this section, we define and characterize an interpretation problem. Informally, an interpretation problem arises from the availability of a set of initial observations from a given system, and of domain knowledge formalized as a set $\mathcal{G} = \{G_1, ..., G_n\}$ of $G^{ap}$ grammars. Every abstraction grammar $G_i \in \mathcal{G}$ generates a set of abstraction patterns that share the same hypothesis $q_h$. The whole set of abstraction patterns that can be generated by $\mathcal{G}$ will be denoted by $\mathcal{P}$.

**Definition 8.** Let $\mathcal{Q}$ be a set of observables and $\mathcal{G}$ a set of abstraction grammars. We say $\mathcal{G}$ induces an **abstraction relation** in $\mathcal{Q} \times \mathcal{Q}$, denoted by $q_i \mid\!\sim q_j$ if and only if there exists an abstraction pattern $P$ generated by some $G \in \mathcal{G}$ such that:

1. $q_j = q_{h_P}$
2. $M_P^{q_i} \cap A_P \neq \varnothing$
3. $q_i \not\mid\!\sim^+ q_i$, where $\mid\!\sim^+$ is the transitive closure of $\mid\!\sim$
4. $image(\mid\!\sim) = \mathcal{Q} - q_0$

The relation $q_i \mid\!\sim q_j$ can be read as *'the observation of $q_i$ allows us to conjecture the presence of $q_j$'*. The transitive closure of the abstraction relation is an strict partial order relation between the domain observables, so that $q_i < q_j \Leftrightarrow q_i \mid\!\sim^+ q_j$, that is, if and only if $\exists q_{k_0}, ..., q_{k_n} \in \mathcal{Q}$ such that $q_{k_0} = q_i$, $q_{k_n} = q_j$ and for all $h$, with $0 \leq h < n$, it holds that $q_{k_h} \mid\!\sim q_{k_{h+1}}$. We denote by $q_i = q_{k_0} \mid\!\sim q_{k_1} \mid\!\sim ... \mid\!\sim q_{k_n} = q_j$ an *abstraction sequence* in n steps that allows to conjecture $q_j$ from $q_i$. This order relation defines an abstraction hierarchy among the observables in $\mathcal{Q}$.

**Example 4.1.** *Let $\mathcal{Q} = \{q_{Pw}, q_{QRS}, q_{Tw}, q_N, q_V, q_{VB}\}$ and $\mathcal{G} = \{G_N, G_{VB}\}$, containing the knowledge represented in examples 3.6 and 3.7. The derived abstraction relation states that $q_{Pw}, q_{QRS}, q_{Tw} \mid\!\sim q_N$, and $q_N, q_V \mid\!\sim q_{VB}$. Intuitively, we can see that this relation splits the observables in the wave, heartbeat, and rhythm levels commonly used in electrocardiogram analysis.*

It is important to note that the abstraction relation is only established between observables in the $A_P$ set. This provides flexibility to define the evidence that form the context of a pattern, as it can belong to the same or even higher abstraction level.

**Definition 9.** We define an **abstraction model** as a tuple $\mathcal{M} = \langle \mathcal{Q}, \mid\!\sim, \mathcal{G} \rangle$, where $\mathcal{Q}$ is the set of domain observables, $\mid\!\sim$ is an abstraction relation between such observables, and $\mathcal{G}$ is the available knowledge as a set of abstraction grammars.

The successive application of the available abstraction grammars result in a series of observations organized in a hierarchy of abstraction, according to the order relation between observables as described above. We are able to define an interpretation problem:

**Definition 10.** We define an **interpretation problem** as a tuple $IP = \langle \mathcal{O}, \mathcal{M} \rangle$, where $\mathcal{O} = (o_1, o_2, \ldots, o_i, \ldots)$ is a sequence of observations requiring interpretation and $\mathcal{M}$ is an abstraction model of the domain.

It is worth mentioning that this definition of an abductive interpretation problem differs from the common definition of abductive diagnosis problem, where the difference between normal and faulty behaviors is explicit, leading to the role of faulty manifestations. Only when some faulty manifestation is detected the abductive process of diagnostic is started. In contrast, in the present framework all the observations have the same status, and the objective of the interpretation process is to provide an interpretation of what is observed at the highest possible abstraction level in terms of the underlying processes. As we will see later, some observables can stand out amongst others regarding the efficiency of the interpretation process, as salient features that can draw some sort of perceptual attention.

As we have said before, any observable $q \in Q_P$ can appear multiple times as different pieces of evidence for an abstraction pattern $P$, in the form of findings collected in the set $M_P$. As a consequence, $P$ can predict multiple observations of the set $\mathcal{O}$ for a given observable $q \in Q_P$, each one of them corresponding to one of the findings of the set $M_P$ through a matching relation. This matching relation is a matter of choice for the agent in charge of the interpretation task, by selecting the observation corresponding to each finding in a given pattern.

**Definition 11.** Given an interpretation problem $IP$, a **matching relation** for a pattern $P \in \mathcal{P}$ is an injective relation in $M_P \times \mathcal{O}$, defined by $m^q \leftarrow o$, if and only if $m^q = \langle \psi, \mathbf{A}, T^b, T^e \rangle \in M_P$ and $o = \langle q, \mathbf{v}, t^b, t^e \rangle \in O(q) \subseteq \mathcal{O}$, such that $\mathbf{A} = \mathbf{v}$, $T^b = t^b$ and $T^e = t^e$.

The assignment included in a matching relation leads us to understand the interpretation problem as a search for a solution for the constraints represented in an abstraction pattern. In this sense, the interpretation problem can be seen as a consistency problem.

From the notion of matching relation we can design a mechanism for abductively interpreting a subset of observations in $\mathcal{O}$ through the use of abstraction patterns. Thus, a matching relation for a given pattern allows us to hypothesize new observations from previous ones, and to iteratively incorporate new evidence to the interpretation by means of a hypothesis generation-and-test cycle. The notion of abstraction hypothesis defines those conditions that a subset of observations must satisfy in order to be abstracted by a new observation, and makes it possible to incrementally build an interpretation by the incorporation of new evidence.

**Definition 12.** Given an interpretation problem $IP$, we define an **abstraction hypothesis** as a tuple $h = \langle o_h, P_h, \leftarrow_h \rangle$, where $P_h = \langle q_h, M_h, C_h, \pi_h \rangle \in \mathcal{P}$, $\leftarrow_h \subseteq M_h \times \mathcal{O}$, and we denote $O_h = image(\leftarrow_h)$, satisfying:

1. $o_h \in O(q_h)$.
2. $o_h = \pi_h(O_h)$.

3. $C_h(\mathbf{A}_h, T_h, \mathbf{A}_1, T_1, ..., \mathbf{A}_n, T_n)$ is consistent.

These conditions entail: (1) an abstraction hypothesis guesses an observation of the observable hypothesized by the pattern; (2) a new observation is obtained from the application of the observation procedure to those observations being assigned to the set of findings $M_h$ by the matching relation; and (3) observations involved in an abstraction hypothesis must satisfy temporal and value consistency.

On the other hand, an abstraction hypothesis assigns a set of observations to the findings of the pattern, giving them the role of evidence for the hypothesis. Even though the matching relation is a matter of choice, and therefore a conjecture by itself, some additional constraints could be considered as default assumptions. An important default assumption in the abstraction of periodic processes states that consecutive observations are related by taking part of the same hypothesis, defining the basic period of the process. This assumption behaves as a sort of operative hypothesis of the abstraction task:

**Default Assumption 2.** (*Basic periodicity) Periodic findings in an abstraction pattern must be assigned consecutive observations by any matching relation:*

$$\forall m_i^q, m_{i+1}^q \in M_h^q, m_i^q \leftarrow_h o_j \wedge \mathrm{q{-}succ}(o_j) \in O_h \Rightarrow m_{i+1}^q \leftarrow_h \mathrm{q{-}succ}(o_j)$$

This default assumption allows us to avoid certain combinations of abstraction hypotheses that, being formally correct, are senseless from an interpretation point of view. For example, without the basic periodicity assumption, a normal rhythm fragment might be abstracted by two alternating bradycardia hypotheses, as shown in figure 6.

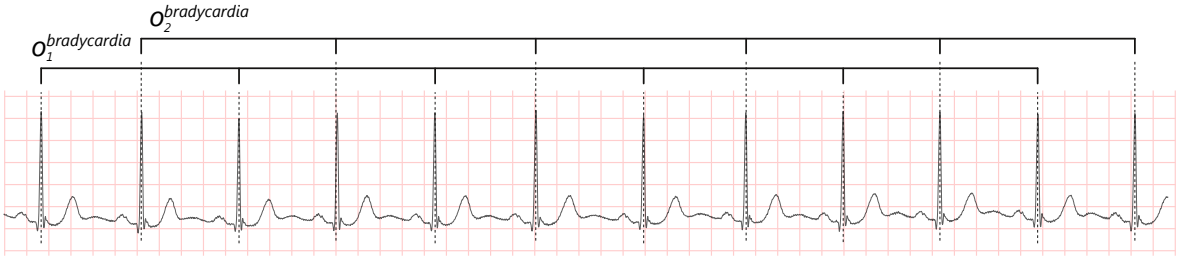


Figure 6: Basic periodicity assumption motivation. [Source: MIT-BIH arrhythmia DB, recording: 103, between 00:40.700 and 00:51.200]

The set of observations that may be abstracted in an interpretation problem $IP$ is $O(domain(\kappa))$, that is, observations corresponding to observables involved in the set of findings to be abstracted by some abstraction pattern. An abstraction hypothesis defines in the set of observations $\mathcal{O}$ a counterpart of the subsets $A_P$ and $E_P$ of the set of findings $M_P$ of a pattern $P$, resulted from selecting a set of observations $O_h \subseteq \mathcal{O}$ by means of a matching relation, satisfying those requirements shown on definition 12.

**Definition 13.** Given an interpretation problem $IP$ and an abstraction hypothesis $h = \langle o_h, P_h, \leftarrow_h \rangle$, we define the following sets of observations:

- $\mathtt{abstracted_by}(o_h) = \{o \in O_h \mid m_i^q \leftarrow_h o \wedge m_i^q \in A_{P_h}\}$.

- $\texttt{environment_of}(o_h) = \{o \in O_h \mid m_i^q \leftarrow_h o \wedge m_i^q \in E_{P_h}\}$.
- $\texttt{evidence_of}(o_h) = \texttt{abstracted_by}(o_h) \cup \texttt{environment_of}(o_h)$.

We denote by $\texttt{abstracted_by}(o_h)$ the set of observations abstracted by $o_h$ and which are somehow its constituents, while $\texttt{environment_of}(o_h)$ denotes the evidential context of $o_h$. We denote by $\texttt{evidence_of}(o_h)$ the set of all observations supporting a specific hypothesis. Since the matching relation is injective, it follows that $\texttt{abstracted_by}(o_h) \cap \texttt{environment_of}(o_h) = \varnothing$.

Definition of these sets can be generalized to include as arguments a set of observations $O = \{o_{h_1}, ..., o_{h_m}\}$ from a set of abstraction hypotheses $h_1, ..., h_m$:

- $\texttt{abstracted_by}(O) = \bigcup_{o_h \in O} \texttt{abstracted_by}(o_h)$
- $\texttt{environment_of}(O) = \bigcup_{o_h \in O} \texttt{environment_of}(o_h)$.
- $\texttt{evidence_of}(O) = \bigcup_{o_h \in O} \texttt{evidence_of}(o_h)$.

As a result of an abstraction hypothesis, a new observation $o_h$ is generated, that can be included in the set of domain observations, so $\mathcal{O} = \mathcal{O} \cup \{o_h\}$. In this way, an interpretation can be incrementally built from the observations, by means of the consistent aggregation of abstraction hypotheses.

**Definition 14.** Given an interpretation problem $IP$, we define an **interpretation** as a consistent set of abstraction hypotheses $I = \{h_1, ..., h_m\}$.

An interpretation can be rewritten as $I = \langle O_I, P_I, \leftarrow_I \rangle$, where: $O_I = \{o_{h_1}, ..., o_{h_m}\}$ is the set of observations guessed by performing multiple abstraction hypotheses; $P_I = \{P_1, ..., P_m\}$ is the set of abstraction patterns used in the interpretation; and $\leftarrow_I = \leftarrow_1 \cup ... \cup \leftarrow_m \subseteq (M_1 \cup ... \cup M_m) \times \mathcal{O}$ is the global matching relation. We should note that the global matching relation $\leftarrow_I$ is not necessarily injective, since some observations could simultaneously belong to $\texttt{abstracted_by}()$ and $\texttt{environment_of}()$ sets of different observations.

From a given interpretation problem $IP$, different interpretations through different sets of abstraction hypothesis can be abductively proposed. Indeed, the definition of interpretation is actually weak, since even an empty set $I = \emptyset$ is formally a valid interpretation. Thus, we need additional criteria in order to select the solution to the interpretation problem as the best choice among different possibilities [24].

**Definition 15.** Given an interpretation problem $IP$, an interpretation $I$ is a **cover** of $IP$ if the set of observations to be interpreted $O(domain(\vdash)) \subseteq \mathcal{O}$ is included in the set of observations abstracted by $I$, that is, $O(domain(\vdash)) \subseteq \texttt{abstracted_by}(O_I)$.

**Definition 16.** Given an interpretation problem $IP$, two different abstraction hypothesis $h$ and $h'$ on mutually exclusive observables $q_h$ and $q_{h'}$ are **alternative hypotheses** if and only if $\texttt{abstracted_by}(o_h) \cap \texttt{abstracted_by}(o_{h'}) \neq \varnothing$.

**Example 4.2.** *A ventricular trigeminy is an infrequent arrhythmia very similar to ventricular bigeminy, except that the ectopic heart beats occur after every pair of normal beats instead of after each one. The grammar to hypothesize a ventricular trigeminy $q_{VT}$ would therefore*

*be very similar to that described in example 3.7, with the difference that each $q_V$ finding would appear after every pair of $q_N$ findings. These two processes are mutually exclusive, in so far as the activation pattern of the heart can be only one in a certain interval. For this reason, in the event of an observation of $q_V$, it may be abstracted by a $q_{VB}$ or a $q_{VT}$ hypothesis, but never by both simultaneously.*

**Definition 17.** Given an interpretation problem $IP$, a cover $I$ for $IP$ is **exclusive** if and only if it contains no alternative hypotheses.

Thus, two or more different hypotheses of mutually exclusive observables abstracted from the same observation will be incompatible in the same interpretation, since inferring a statement and its negation is logically prevented, and therefore only one of them can be selected.

On the other hand, a parsimony criterion is required, in order to disambiguate among possible interpretations to select the most plausible ones as those whose complexity is minimum [24]. We translate this minimum complexity in terms of minimal cardinality.

**Definition 18.** Given an interpretation problem $IP$, a cover $I$ for $IP$ is **minimal**, if and only if its cardinality is smallest among all covers for $IP$.

Minimality introduces a parsimony criterion on hypothesis generation, promoting temporally maximal hypotheses, that is, those hypothesis of a larger scope instead of multiple equivalent hypotheses of smaller scope. For example, consider an abstraction pattern that allows to conjecture regular cardiac rhythm from the presence of three or more consecutive heart beats. Without a parsimony criterion, a sequence of nine consecutive beats could be abstracted by up to three consecutive rhythm observations, even when a single rhythm observation would be sufficient and better.

**Definition 19.** The **solution** of an interpretation problem $IP$, is the set of all minimal and exclusive covers of $IP$.

This definition of solution is very conservative and with a limited practical value, since the usual objective is to obtain a small set of interpretations explaining what has been observed (ideally only one). However, it allows us to characterize the problem in terms of complexity. Abduction has been formulated under different frameworks according to the task to be addressed, but it has always shown as an intractable problem in the general case. The next theorem proves that an interpretation problem is also an intractable problem.

**Theorem 1.** *Finding the solution to an interpretation problem is NP-hard.*

*Proof:* We will provide a polynomial-time reduction of the well-known set covering problem to an interpretation problem. Given a set of elements $U = \{u_1, ..., u_m\}$ and a set $S$ of subsets of $U$, a cover is a set $C \subseteq S$ of subsets of $S$ whose union is $U$. Regarding complexity analysis two different problems of interest are identified:

- A set covering decision problem, stating that given a pair $(U, S)$ and an integer $k$ the question is whether there is a set covering of size $k$ or less. This decision version of set covering is NP-complete.

- A set covering optimization problem, stating that given a pair $(U, S)$ the task is to find a set covering that uses the fewest sets. This optimization version of set covering is NP-hard.

We will therefore reduce the set covering problem to an interpretation problem by means of a polynomial-time function $\varphi$. Thus, we shall prove that $\varphi(U, S)$ is an interpretation problem, and there is a set covering of $\varphi(U, S)$ of size $k$ or less if and only if there is a set covering of $U$ in $S$ of size $k$ or less.

Given a pair $(U, S)$, let $\varphi(U, S) = \langle \mathcal{O}, \mathcal{M} \rangle$ where:

1. $\mathcal{O} = U = \{u_1, ..., u_m\}$, such that $u_i = \langle q, true, i \rangle$ and $q = \langle \psi, present, T \rangle$.
2. $\mathcal{M} = \langle Q, \Vdash, \mathcal{P} \rangle$, such that $domain(\Vdash) = q$.
3. $\forall s = \{u_{i_1}, ..., u_{i_n}\} \in S,\ \exists P \in \mathcal{P}$, being $P = \langle q_P, M_P, C_P, \pi_P \rangle$, where:
   - $q \Vdash q_P$ and $P \neq P' \Rightarrow q_P \neq q_{P'}$.
   - $M_P = A_P = M_P^q = \{m_1^q = \langle \psi, present, T_1 \rangle, ..., m_n^q\}$.
   - $C_P = u_1.present \wedge ... \wedge u_n.present = true$.
   - $\pi_P = \bigwedge_{k=1}^{n} (u_k.t = i_k)$.

Thus, $\varphi(U, S)$ is an interpretation problem according to this definition. On the other hand, $\varphi(U, S)$ can be built in polynomial time. In addition, for all $s \in S$ there exists an abstraction hypothesis $h = \langle o_h, P_h, \leftarrow_h \rangle$ such that:

1. $o_h = \langle q_{P_h}, \pi_{P_h} = true, min_{u_i \in s}\{i\}, max_{u_i \in s}\{i\} \rangle$.
2. $u_i \in s \Rightarrow u_i \in image(\leftarrow_h)$.
3. $\leftarrow_h$ is consistent, since the set of observations satisfying $\pi_{P_h} = true$ also satisfies the constraints in $C_P$.

Since any abstraction hypothesis involves a different abstraction pattern there are no alternative hypothesis in any interpretation of $\varphi(U, S)$. In addition, any combination of temporal constraints from abstraction patterns in $\varphi(U, S)$ is, by definition, consistent.

Suppose there is a set covering $C \subseteq S$ of $U$ of size $k$ or less. For all $u \in U$ there exists $c_i \in C - \{\varnothing\}$ such that $u \in c_i$ and, by the above construction, there exists $h_i \in I$ such that $\texttt{abstracted_by}(o_{h_i}) = \{u \in image(\leftarrow_{h_i})\} = \{u \in c_i\} = c_i$, and therefore, $O(domain(\Vdash)) \subseteq \bigcup_{h_i \in I} \texttt{abstracted_by}(o_{h_i}) = \bigcup_i c_i = C$, that is, the set of abstraction hypotheses $I$ is an exclusive cover of the interpretation problem $\varphi(U, S)$ of size $k$ or less.

Following the same reasoning as for the set covering optimization problem, finding a minimal and a exclusive cover of an interpretation problem $\varphi(U, S)$ is NP-hard, since we can use the solution of this problem to check if there is an exclusive cover of the interpretation problem of size $k$ or less, and this has just been proved to be NP-complete. □

# 5 Solving an interpretation problem: A heuristic search approach

The solution set for an interpretation problem $IP$ consists of all exclusive covers of $IP$ having the minimum possible number of abstraction hypotheses. Obtaining this solution set can be stated as a search on the set of interpretations of $IP$. The major source of complexity of searching for a solution is the local election, from the available evidence in $\mathcal{O}$, of the most appropriate matching relation for a number of abstraction hypotheses that can globally shape a minimal and exclusive cover of $IP$.

Nevertheless, the whole notion of solution must be revised in practical terms, due to the intractability of the task and the incompleteness of the abstraction model, that is, of the available knowledge. Indeed, we assume that any realistic abstraction model can hardly provide a cover for every possible interpretation problem. Hence the objective should move from searching for a solution to searching for an approximate solution.

Certain principles applicable to the interpretation problem can be exploited in order to approach a solution in an iterative way, bounding the combinatorial complexity of the search. These principles can be stated as a set of heuristics that make it possible to evaluate and discriminate some interpretations against others from the same base evidence:

- A *coverage principle*, which states the preference for interpretations explaining more initial observations.
- A *simplicity principle*, which states the preference for interpretations with fewer abstraction hypotheses.
- An *abstraction principle*, which states the preference for interpretations involving higher abstraction levels.
- A *predictability principle*, which states the preference for interpretations that properly predict future evidence.

The coverage and simplicity principles are used to define a cost measure for the heuristic search process [9], while the abstraction and predictability principles are used to guide the reasoning process, trying to emulate the same shortcuts used by humans.

Given an interpretation problem $IP$, a heuristic vector for a certain interpretation $I$ can be defined to guide the search, as $\epsilon(I) = (1 - \varsigma(I), \kappa(I))$, where $\varsigma(I)$ is the *covering ratio* of $I$, defined as $\varsigma(I) = |\texttt{abstracted_by}(O_I)|/|O(domain(\mathbb{k}))|$, and $\kappa(I) = |O_I|$ is the *complexity* of $I$. The main goal of the search strategy is to approach a solution with a maximum covering ratio and a minimum complexity, which is equivalent to the minimization of the heuristic vector. Covering ratio will be considered the primary heuristic, and complexity will be considered for ranking those interpretations with the same covering ratio. The $\epsilon(I)$ heuristic is intuitive and very easy to calculate, but as a counterpart it is a non-admissible heuristic, since it is not monotone and can alternatively underestimate and overestimate the true goal covering. Therefore optimality cannot be guaranteed and we require an algorithm efficient with this type of heuristic. We propose the CONSTRUE() algorithm, whose pseudocode is shown in algorithm 1. This algorithm is a small variation of the K-Best First Search algorithm [10], with partial expansion to reduce the number of explored nodes.

**Algorithm 1** CONSTRUE search algorithm.

```
 1: function CONSTRUE(IP)
 2:     var I_0 = ∅
 3:     var K = max(|{q_j ∈ Q | q_i ⊰ q_j, q_i ∈ Q}|)
 4:     SET_FOCUS(I_0, o_1)
 5:     var open = SORTED([⟨ϵ(I_0), I_0⟩])
 6:     var closed = SORTED([])
 7:     while open ≠ ∅ do
 8:         for all I ∈ open[0...K] do
 9:             I' = NEXT(GET_DESCENDANTS(I))
10:             if I' is null then
11:                 open = open − {⟨ϵ(I), I⟩}
12:                 closed = closed ∪ {⟨ϵ(I), I⟩}
13:             else if ς(I') = 1.0 then
14:                 return I'
15:             else
16:                 open = open ∪ {⟨ϵ(I'), I'⟩}
17:             end if
18:         end for
19:     end while
20:     return min(closed)
21: end function
```

The CONSTRUE() algorithm takes as input an interpretation problem $IP$, and returns the first found interpretation with full coverage, or the interpretation with the maximum covering ratio and minimum complexity if no covers are found, using the abstraction and predictability principles along the searching process as well. For this, it manages two ordered lists of interpretations, named $open$ and $closed$. Every interpretation is annotated with the computed values of the heuristic vector. The $open$ list contains those partial interpretations that can further evolve by (1) appending new hypotheses or (2) extending previously conjectured hypotheses to subsume or predict new evidence. This $open$ list is initialized with the trivial interpretation $I_0 = \emptyset$. The $closed$ list contains those interpretations that cannot explain more evidence.

At each iteration, the algorithm selects the $K$ most promising interpretations according to the heuristic vector (line 8), and partially expands each one of them to obtain the next descendant node $I'$. If this node is a solution, then the process finishes by returning it (line 13), otherwise it is added to the $open$ list. The partial expansion warrants that the $open$ list grows at each iteration at most $K$ new nodes, in order to save memory. When a node cannot expand more, it is added to the $closed$ list (line 12), from which the solution is taken if no full coverages are found (line 20).

The selection of a value for the $K$ parameter depends on the problem at hand. We select its value as $K = max(|\{q_j \in \mathcal{Q} \mid q_i \lhd q_j, q_i \in \mathcal{Q}\}|)$, that is, as the maximum number of observables that can be abstracted from any observable $q_i$. The intuition behind this choice is that at any point in the interpretation process, and with the same heuristic values, we will give the same chance to any plausible abstraction hypothesis in order to explain a certain observation.

In order to expand the current set of interpretations, the GET_DESCENDANTS() function relies on different reasoning modes, that is, different forms of abduction and deduction, which are brought into play under the guidance of an attentional mechanism. Since searching for a solution involves at the end the election of a matching relation, both observations and findings should be included in the scope of this mechanism. Hence, a focus of attention can be defined to answer the following question: which is the next observation or finding to be processed? Answering this question develops a hypothesis-and-test cycle: if the attention focuses on an observation, then an abstraction hypothesis explaining such observation should be generated (hypothesize); but if the attention focuses on a finding predicted by some hypothesis, an observation should be sought to match such finding (test).

To illustrate and motivate the reasoning modes implemented to build interpretations and support the execution of the CONSTRUE() algorithm, we use a simple, but complete, interpretation problem.

**Example 5.1.** *Let $\mathcal{Q} = \{q_{wave}, q_{Pw}, q_{QRS}, q_{Tw}, q_N\}, \mathcal{G} = \{G_w, G_N, G_{Tw}\}$, where $G_w$ models the example 3.5, $G_N$ is described in example 3.6, and $G_{Tw}$ describes the knowledge to conjecture a T wave as follows:*

$$G_{Tw} = ((\{D\}, \{q_{QRS}, q_{wave}\}, \{H = q_{Tw} \rightarrow q_{QRS}[c_1]D \\ D \rightarrow q_{wave}[c_2]\}), \emptyset, \texttt{Tw_delin}())$$

*where the set of constraints is defined as following: $c_{11} \equiv \langle 80ms \leq T^b_{Tw} - T^e_{QRS} \leq 120ms\rangle$, $c_{12} \equiv \langle T^e_{Tw} - T^b_{QRS} \leq 520ms\rangle$, $c_{21} \equiv \langle T^b_{Tw} = T^b_{wave}\rangle$, $c_{22} \equiv \langle T^e_{Tw} = T^e_{wave}\rangle$ and $c_{23} \equiv \langle max(diff(sig[m^{wave}]) \leq 0.7 \cdot max(diff(sig[m^{QRS}]))\rangle$.*

*This grammar hypothesizes the observation of a T wave shortly after the observation of a QRS complex, requiring a significant decrease in the maximum slope of the signal. Let us note that the observation of a T wave is not deduced from the observation of a QRS complex since it is not a necessary condition for the observation of a QRS complex. Attribute computation rules are denoted as `Tw_delin()`, and represent a procedure to obtain the time limits of a T wave based on the observation of a leading QRS complex. This method may be any of those described in the literature, such as [19].*

*In addition to the $P_{wave}$ pattern generated by $G_w$ and detailed in example 3.5, $G_N$ and $G_{Tw}$ generate the following abstraction patterns:*

$$P_N = \langle q_N, A_{P_N} = \{m^{Pw}, m^{QRS}, m^{Tw}\} \cup E_{P_N} = \emptyset, C_{P_N}, \emptyset\rangle$$
$$P_{Tw} = \langle q_{Tw}, A_{P_{Tw}} = \{m^{wave}\} \cup E_{P_{Tw}} = \{m^{QRS}\}, c_1 \cup c_2, \texttt{Tw_delin}()\rangle$$

*Finally, let $\mathcal{O} = \{o^{wave}_1 = \langle q_{wave}, \emptyset, 0.300, 0.403\rangle, o^{wave}_2 = \langle q_{wave}, \emptyset, 0.463, 0.549\rangle, o^{Pw} = \langle q_{Pw}, \emptyset, 0.300, 0.403\rangle, o^{QRS} = \langle q_{QRS}, \emptyset, 0.463, 0.549\rangle\}$ be a set of initial observations including a P wave and a QRS complex abstracting two wave observations located in specific time points.*

*Given this interpretation problem, figure 7 shows the starting point for the interpretation: where the root of the interpretation process is the trivial interpretation $I_0$, and the attention is focused on the first observation. The sequence of reasoning steps towards the resolution of this interpretation problem will be explained in the next subsections.*

## 5.1 Focus of attention

The focus of attention is modeled as a stack, so once the focus is set on a particular observation (or finding), any observation that was previously under the focus will not return to be focused on until the reasoning process on the current observation is finished. Algorithm 2 shows how the different reasoning modes are invoked based on the content of the focus of attention, resulting in a hypothesize-and-test cycle.

**Algorithm 2** Method for obtaining the descendants of an interpretation using different reasoning modes based on the content of the focus of attention.

```
1:  function GET_DESCENDANTS(I)
2:      var focus = GET_FOCUS(I).TOP()
3:      var desc = ∅
4:      if IS_OBSERVATION(focus) then
5:          if focus = o_h | h ∈ I then
6:              desc = DEDUCE(I, focus)
7:          end if
8:          desc = desc ∪ ABDUCE(I, focus) ∪ ADVANCE(I, focus)
9:      else if IS_FINDING(focus) then
10:         desc = SUBSUME(I, focus) ∪ PREDICT(I, focus)
11:     end if
12:     return desc
13: end function
```

Lines 4-8 generate the descendants of an interpretation $I$ when at the top of the stack there is an observation. These descendants are the result of two possible reasoning modes: the deduction of new findings, performed by the DEDUCE() function provided that the observation being focused on is an abstraction hypothesis; and the abduction of a new hypothesis explaining the observation being focused on, performed by the ABDUCE() function. A last descendant is obtained using the ADVANCE() function, which simply restores the previous focus of attention by means of a POP() operation. If the focus is then empty, ADVANCE() inserts the following observation to explain, that may be selected by temporal order in the general case, or by some domain-dependent saliency criterion to prioritize certain observations over others. By removing an observation at the top of the focus of attention, the ADVANCE() function sets aside that observation as unintelligible to the current interpretation, according to the available knowledge.

If the top of the stack contains a finding, then algorithm 2 obtains the descendants of the interpretation from the SUBSUME() and PREDICT() functions (line 10). The first one looks for an existing observation satisfying the constraints on the finding being focused on, while the second makes predictions about observables that have not yet been observed. Below we detail all these reasoning modes separately, and we will illustrate how the CONSTRUE() algorithm combines them in order to solve the interpretation problem described in example 5.1.

## 5.2 Building an interpretation: Abduction

Algorithm 3 enables the abductive generation of new abstraction hypotheses. It is applied when the attention is focused on an observation that can be abstracted by some abstraction pattern, producing a new observation at a higher level of abstraction.

**Algorithm 3** Moving forward an interpretation through abduction.

```
1: function ABDUCE(I, o_i)
2:     var desc = ∅
3:     for all G = ⟨V_N, V_T, q_H, R⟩ ∈ 𝒢 | q(o_i) ⊧ q_H do
4:         for all (D → q[c]F) ∈ R | q(o_i) is_a q ∧ [c] → m^q ∈ A_P do
5:             P_h ← ⟨q_H, {m^q}, [c], π_h⟩
6:             h = ⟨o_h, P_h, {m^q ← o_i}⟩
7:             B_h = D; E_h = F
8:             I' = ⟨O_I ∪ {o_h}, P_I ∪ {P_h}, ←_I ∪ ←_h⟩
9:             𝒪 = 𝒪 ∪ {o_h}
10:            GET_FOCUS(I').POP()
11:            GET_FOCUS(I').PUSH(o_h)
12:            desc = desc ∪ {I'}
13:        end for
14:    end for
15:    return desc
16: end function
```

The result of ABDUCE() is a set of interpretations $I'$, each one adding a new abstraction hypothesis with respect to the parent interpretation $I$. To generate these abstraction hypotheses, we iterate through those grammars that can make a conjecture from the observation $o_i$ being focused on (line 3). Then, for each grammar, each rule including the corresponding observable $q(o_i)$ (line 4) makes use of the above mentioned constructive method to initialize an abstraction pattern with a single finding of this observable (line 5), and a new hypothesis is conjectured with a matching relation involving both the observation being focused on and the finding (line 6). Two additional variables, $B_h$ and $E_h$, allow us to mark the involved rule and they will play an important role in subsequent steps (line 7). Then, the set of constraints concerning this finding are checked, and finally the new hypothesis opens a new interpretation (lines 8-9).

In this way, the ABDUCE() function implements, from a single piece of evidence, the hypothesize step of the hypothesize-and-test cycle. Below we explain the reasoning modes involved in the test step of the cycle.

**Example 5.2.** *Consider the interpretation problem set out in example 5.1 and the interpretation $I_0$ shown in figure 7. According to algorithm 2, the ABDUCE() function is used to move forward the interpretation, since the focus of attention points to an observation $o^{Pw}$. The abstraction pattern that supports this operation is $P_N$, and a matching relation is established with the $m^{Pw}$ finding. As a result, the following hypothesis is generated:*

$$h_1 = \langle o^N, P_N, \{m^{Pw} \leftarrow o^{Pw}\}\rangle$$

*Figure 7 shows the result of this reasoning process, in a new interpretation called $I_1$. Note that the focus of attention has been moved to the newly created hypothesis (lines 10-11 of the ABDUCE() function).*

## 5.3 Building an interpretation: Deduction

This reasoning mode is applied when the attention is focused on an observation $o_h$ previously conjectured as part of an abstraction hypothesis $h = \langle o_h, P_h, \leftarrow_h \rangle$ (see algorithm 4). The DEDUCE() function takes the evidence that has led to conjecture $o_h$ and tries to extend it with new findings that can be expected, that is, deduced, from the abstraction grammar used to guess the observation. The key point is that this deduction process follows an iterative procedure, as the corresponding abstraction pattern is dynamically generated from the grammar. Hence the DEDUCE() function aims to extend a partial matching relation by providing the next finding to be tested, as part of the test step of the hypothesize-and-test cycle.

**Algorithm 4** Moving forward an interpretation through deduction of new findings.

1: **function** DEDUCE($I, o_h$)
2: **var** $desc = \emptyset$
3: **if** $B_h \neq H$ **then**
4: **for all** $(X \rightarrow q'[c]B_h) \in R \mid X \in V_n$ **do**
5: $P'_h \leftarrow \langle q_h, M_h \cup \{m_0^{q'}\}, C_h \cup C_c, \pi_h \rangle$
6: $B_h = X$
7: $I' = \langle O_I, P_I \cup \{P'_h\}, \leftarrow_I \rangle$
8: GET_FOCUS($I'$).PUSH($m_0^{q'}$)
9: $desc = desc \cup \{I'\}$
10: **end for**
11: **else**
12: **for all** $(E_h \rightarrow q'[c]X) \in R \mid X \in V_n$ **do**
13: $P'_h \leftarrow \langle q_h, M_h \cup \{m_{n+1}^{q'}\}, C_h \cup C_c, \pi_h \rangle$
14: $E_h = X$
15: $I' = \langle O_I, P_I \cup \{P'_h\}, \leftarrow_I \rangle$
16: GET_FOCUS($I'$).PUSH($m_{n+1}^{q'}$)
17: $desc = desc \cup \{I'\}$
18: **end for**
19: **end if**
20: **return** $desc$
21: **end function**

Since the first finding that has led to conjecture $o_h$ does not have to appear at the beginning of the grammar description, the corresponding abstraction pattern will not be, in general, generated incrementally from the first rule of the grammar. Taking as a starting point the rule used to conjecture $o_h$ (line 4 in algorithm 3), the idea is to add a new finding by adding a new rule at both sides, towards the beginning and the end of the grammar, using the $B_h$(egin) and $E_h$(nd) variables in order to track the next finding to be tested. In this way, the matching relation can be incrementally built from a sequence of findings. $B_h$ represents the non-terminal at the left side of the rule that generated the first finding in the matching relation, while $E_h$ represents the non-terminal at the right side of the rule that generated the last finding in the matching relation. With this information we are able to extend the current abstraction pattern $P_h$ in two opposite directions:

- Towards the beginning of the grammar (lines 3-10). We explore the set of observables that may occur before the first finding according to the rules of the grammar (line 4), and we deduce a new finding for each one of them in different descendant interpretations. An update of the current abstraction pattern is performed (line 5), and, consequently, an update of the current interpretation (line 7). The newly deduced finding is the next to be focused on (line 8).
- Towards the end of the grammar (lines 11-18). For each one of the observables that may occur after the last finding, a new finding is deduced, updating the abstraction pattern and the interpretation. This new finding is the next to be focused on (lines 12-18).

**Example 5.3.** *Consider the interpretation problem set out in example 5.1 and the interpretation $I_1$ shown in figure 7. Remember that the grammar used to generate the hypothesis in the focus of attention, $G_N$, has the following form:*

$$\begin{aligned} H = q_N &\rightarrow q_{Pw}[c_1]D \\ D &\rightarrow q_{QRS}[c_2]E \\ E &\rightarrow q_{Tw}[c_3] \end{aligned}$$

*In this situation, it is possible to deduce new findings from the $o^N$ hypothesis. Following algorithm 3 we can check that $B_h = H$ and $E_h = D$, since the only finding in the matching relation is $m^{Pw}$. Deduction has then to be performed after this last finding, using the rule $D \rightarrow q_{QRS}[c_2]E$. After $[c_2]$ constraint checking, the resulting finding is the following:*

$$m_{n+1}^{q'} = m^{QRS} = \langle q_{QRS}, \emptyset, T^b_{QRS} \in [0.400, 0.520], T^e_{QRS} \in [0.450, 0.660] \rangle$$

*Figure 7 illustrates the outcome of this reasoning process and the uncertainty in the temporal limits of the predicted finding, that is now being focused on in the interpretation $I_2$.*

## 5.4 Building an interpretation: Subsumption

Subsumption is performed when the attention is focused on a finding previously deduced from some abstraction grammar (see algorithm 5). This reasoning mode avoids to generate a new hypothesis for every piece of available evidence if it can be explained by a previous hypothesis. The SUBSUME() function explores the set of observations $\mathcal{O}$ and selects those consistent with the constraints of the finding in the focus of attention (line 3), expanding the matching relation of the corresponding hypothesis in different descendant interpretations (line 4). The focus of attention is then restored to its previous state (line 5), allowing to deduce new findings from the same hypothesis. The SUBSUME() function clearly enforces the simplicity principle.

**Example 5.4.** *Consider the interpretation $I_2$ shown in figure 7. If we apply the subsumption procedure, it is possible to set a matching relation between $o^{QRS}$ and $m^{QRS}$, since this observation satisfies all the constraints on the finding. The result is shown in the interpretation $I_3$. Note that the uncertainty in the end time of the $o^N$ hypothesis is now reduced after the matching, having $T^e_N \in [0.631, 1.030]$. Then the attention focuses once again on this hypothesis, and therefore a new deduction operation may be performed.*

**Algorithm 5** Moving forward an interpretation through subsumption.

1: **function** SUBSUME($I, m_i$)
2: **var** $desc = \emptyset$
3: **for all** $o_j \in \mathcal{O} \mid m_i \leftarrow o_j$ **do**
4: $I' = \langle O_I, P_I, \leftarrow_I \cup \{m_i \leftarrow o_j\}\rangle$
5: GET_FOCUS($I'$).POP($m_i$)
6: $desc = desc \cup \{I'\}$
7: **end for**
8: **return** $desc$
9: **end function**

## 5.5 Building an interpretation: Prediction

This reasoning mode is also performed when the attention is focused on a finding previously deduced from some abstraction grammar (see algorithm 6). In this case, if a finding previously deduced has not been observed yet it will be predicted.

**Algorithm 6** Moving forward an interpretation through prediction of non-available evidence.

1: **function** PREDICT($I, m_i$)
2: **var** $desc = \emptyset$
3: **for all** $\langle V_N, V_T, q_H, R\rangle \in \mathcal{G} \mid q_H$ `is_a` $q(m_i)$ **do**
4: $P_h \leftarrow \langle q_H, \emptyset, C_h, \pi_h\rangle$
5: $h = \langle o_h, P_h, \emptyset\rangle$
6: $B_h = E_h = H$
7: $I' = \langle O_I \cup \{o_h\}, P_I \cup \{P_h\}, \leftarrow_I \cup \{m_i \leftarrow o_h\}\rangle$
8: $\mathcal{O} = \mathcal{O} \cup \{o_h\}$
9: GET_FOCUS($I'$).POP($m_i$)
10: GET_FOCUS($I'$).PUSH($o_h$)
11: $desc = desc \cup \{I'\}$
12: **end for**
13: **return** $desc$
14: **end function**

The goal of the PREDICT() function is to conjecture a new observation for the finding being focused on. To do this, the knowledge base is explored and those grammars whose hypothesized observable is more specific than the predicted observable are selected (line 3). Then, a new pattern is initialized with no evidence supporting it, and a new abstraction hypothesis with an empty matching relation is generated (lines 4-5). Finally, the attention focuses on the observation being guessed (lines 9-10) to enable the DEDUCE() function to start a new iteration of the test step at a lower abstraction level. Since the $B_h$ and $E_h$ variables point to the first non-terminal symbol of the grammar (line 6), the corresponding abstraction pattern is generated only towards the end of the grammar.

**Example 5.5.** *Starting from the $I_3$ interpretation shown in figure 7, the next step we can take to move forward the interpretation is a new deduction on the $o^N$ hypothesis, generating a new finding $m^{Tw}$, leading to the $I_4$ interpretation. Since there is no available observation of the T wave, a matching with this new finding $m^{Tw}$ cannot be made by the SUBSUME() function, so the only option to move forward this interpretation is through*

*prediction. Following the PREDICT() function, the $G_{Tw}$ grammar can be selected, and a new observation $o^{Tw}$ can be conjectured, generating the $I_5$ interpretation.*

*From $I_5$ we can continue the deduction on the $o^{Tw}$ hypothesis. If we apply the DEDUCE() function we obtain the $m^{QRS'}$ finding from the environment, shown in figure 7 as $I_6$. To move on, we can apply the SUBSUME() function, establishing the matching relation $\{m^{QRS'} \leftarrow o^{QRS}\}$. This leads to the $I_7$ interpretation, in which the uncertainty on the $o^{Tw}$ observation is reduced, but the evidence for the $P_{Tw}$ pattern is not yet complete. A new DEDUCE() step is necessary, which deduces the $m^{wave}$ necessary finding in the $I_8$ interpretation. This finding is also absent, so another PREDICT() step is required. In this last step, the $P_{wave}$ pattern can be applied to observe the deviation in the raw ECG signal, generating the $o_3^{wave}$ observation and completing the necessary evidence for the $o^{Tw}$ observation and thus also for $o^N$. Constraint solving assigns the value of $t^b_{Tw}$, $t^e_{Tw}$ and $t^e_N$, so the result is a cover of the initial interpretation problem in which all the hypotheses have a necessary and sufficient set of evidence. This solution is depicted in $I_9$.*

*It is worth noting that in this example the global matching relation $\leftarrow_I$ is not injective, since $m^{QRS} \leftarrow o^{QRS}$ and $m^{QRS'} \leftarrow o^{QRS}$. Also note that each interpretation only generates one descendant, but in a more complex scenario the possibilities are numerous, and the responsibility to find the proper sequence of reasoning steps lies with the CONSTRUE() algorithm.*

## 5.6 Improving the efficiency of interpretation through saliency

Starting a hypothesize-and-test cycle for every single sample is not feasible for most of the time series interpretation problems. Still, many problems may benefit from certain saliency features that can guide the attention focus to some limited temporal fragments that can be easily interpretable. Thus, the interpretation of the whole time series can pivot on a reduced number of initial observations, therefore speeding-up the interpretation process.

A saliency-based attentional strategy can be devised from the definition of abstraction patterns, by using a subset of their constraints as a coarse filter to identify a set of plausible observations. For example, in the ECG interpretation problem, the minimum slope constraint defined by the QRS abstraction pattern is a reliable feature to select more promising signal segments to focus the interpretation on, in the same way that a cardiologist focuses on the prominent peaks in the signal to start the analysis [31].

By using some of the constraints of the abstraction patterns we are taking advantage of the necessary nature of these constraints. If there are multiple alternative constraints in the description of an abstraction pattern, those that are easier to compute can be preferable, even when some of the potential observations could be lost, as they can be predicted later as part of the hypothesize-and-test cycle. This is frequent when a cardiologist makes a fast read of an ECG. Additionally, the non-sufficient nature of these constraints makes it necessary to insert the resulting set of plausible observations in the hypothesize-and-test cycle where they can be refuted later.

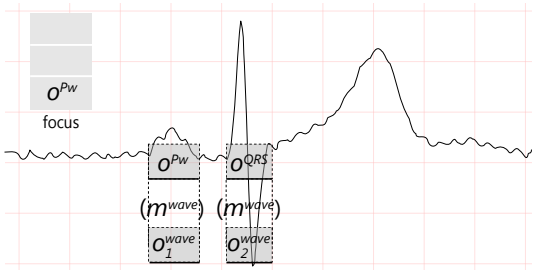


$I_0$ = Initial evidence

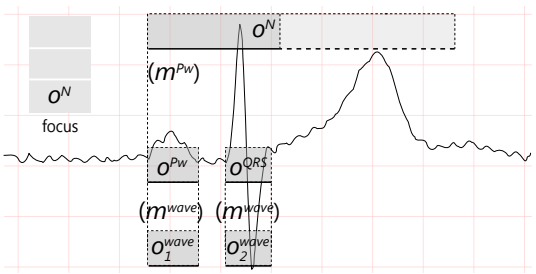


$I_1$ = ABDUCE($I_0, o^{Pw}$)

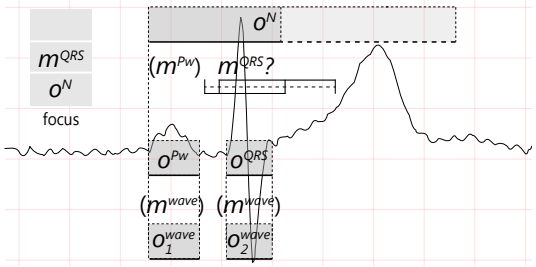


$I_2$ = DEDUCE($I_1, o^N$)

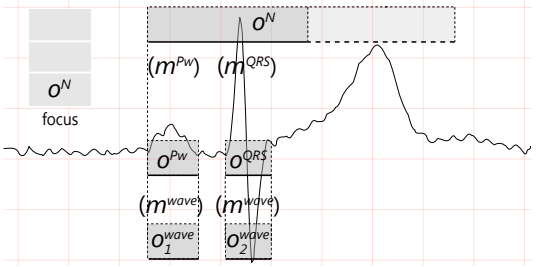


$I_3$ = SUBSUME($I_2, m^{QRS}$)

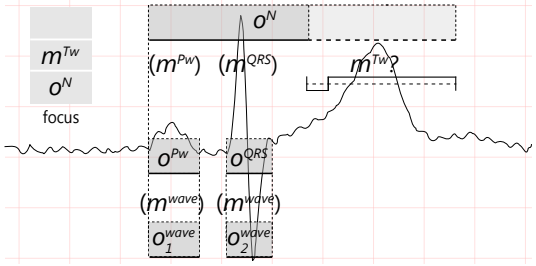


$I_4$ = DEDUCE($I_3, o^N$)

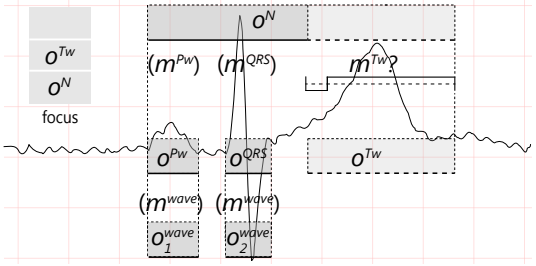


$I_5$ = PREDICT($I_4, m^{Tw}$)

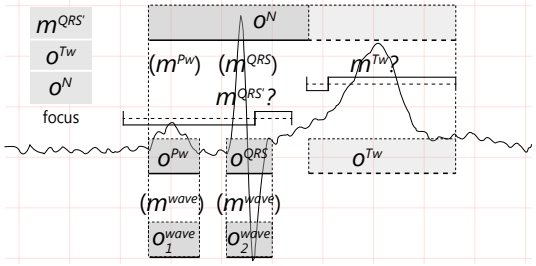


$I_6$ = DEDUCE($I_5, o^{Tw}$)

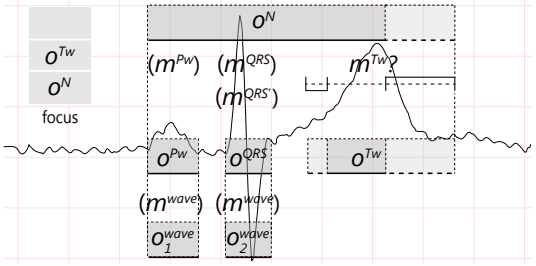


$I_7$ = SUBSUME($I_6, m^{QRS'}$)

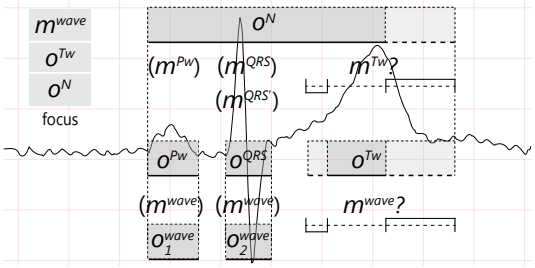


$I_8$ = DEDUCE($I_7, o^{Tw}$)

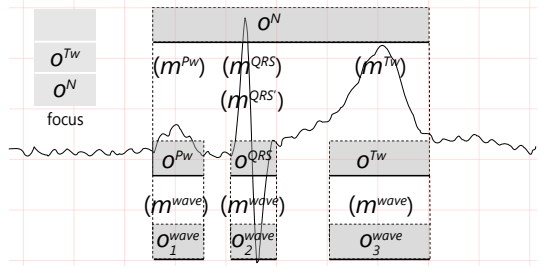


$I_9$ = PREDICT($I_8, m^{wave}$)

Figure 7: Sequence of reasoning steps to solve a simple interpretation problem.

# 6 Strengths of the model: Overcoming some limitations of classifiers

In this section we provide some practical examples to illustrate some of the strengths of the proposed interpretation framework, and its ability to tackle with typical weaknesses of classification-based approaches.

## 6.1 Avoiding a casuistry-based interpretation

In the time domain, a classification-based recognition of multiple processes occurring concurrently usually lead to a casuistry-based proliferation of classes, where it is usually needed a new class for each possible superposition of processes, in order to properly identify all situations. It is common to use a representation in transform domain, where certain regular processes are easily separable, though at the expense of a cumbersome representation of the temporal information [21]. In contrast, in the present framework, the hypothesize-and-test cycle aims to conjecture those hypotheses that best explain the available evidence, including in a natural way simultaneous hypotheses as long as they are consistent and they are not mutually exclusive.

ECG interpretation provides some interesting examples of this sort of problems. Atrial fibrillation, a common heart arrhythmia caused by the independent and erratic contractions of the atrial muscle fibers, is characterized by an irregularly irregular heart rhythm [31]. Most of the classification techniques for the identification of atrial fibrillation are based on the analysis of the time interval between consecutive beats, trying to detect this irregularity [25]. These techniques offer good results in those situations in which atrial fibrillation is the only anomaly, but they fail to properly identify complex scenarios which go beyond the distinction between atrial fibrillation and normal rhythm. In the strip of figure 8, obtained during a pilot study for the home follow-up of patients with cardiac diseases [28], such a classifier would wrongly identify this segment as an atrial fibrillation episode, since the observed rhythm variability is consistent with the description of this arrhythmia. In contrast, the present interpretation framework properly explains the first 5 beats as a sinus bradycardia, compatible with the presence of a premature ectopic beat in the second position, followed by a trigeminy pattern during 6 beats, and finally another ectopic beat with a morphology change. The reason to choose this interpretation, despite being more complex than the atrial fibrillation explanation, is that it is able to abstract some of the small P waves before the QRS complexes, increasing the interpretation coverage.

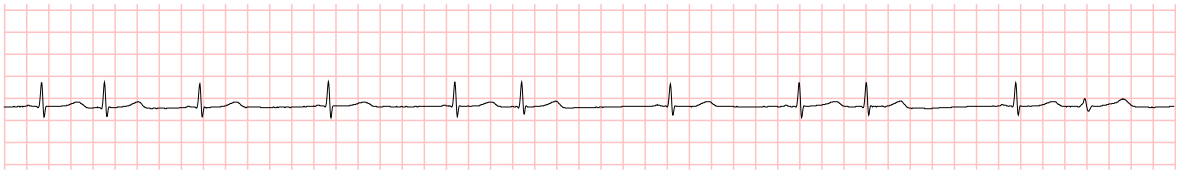

Figure 8: False atrial fibrillation episode. [Source: Mobiguide Project [28], private recording]

## 6.2 Coping with ignorance

Most of the classifiers solve a separability problem among classes, whether by learning from a training set or by eliciting prior knowledge, and they are implicitly based on the closed-world assumption, i.e., every new instance to be classified is assigned to one of the predefined classes. Such classifiers may additionally include a reject option for all those instances that could be misclassified as they appear too close to the classification boundaries [4, 13]. This reject option is added as another possible answer expressing doubt. However, such classifiers fail to classify new instances of unknown classes, since they cannot express ignorance. An approach to this problem can be found in novelty detection proposals [26], which can detect when a new instance does not fit any of the predefined classes as it substantially differs from those instances available during training. Still, they are limited to a common feature representation for every instance, hindering to identify what is unintelligible from the available knowledge.

The present framework provides an expression of ignorance as a common result of the interpretation problem. As long as the abstraction model may be incomplete, a non coverage of some piece of evidence by any interpretation is an expression of partial ignorance. In the extreme case, the trivial interpretation $I_0$ can be a proper solution for an interpretation problem, expressing total ignorance. Furthermore, abduction naturally includes the notion of ignorance in the reasoning process, since any single piece of evidence can be sufficient to guess an interpretation, and the hypothesize-and-test cycle can be understood as a process of incremental addition of evidence against an initial state of ignorance, being able to provide an interpretation at any time based on the available evidence.

As an example, consider the interpretation problem illustrated in figure 9. Let the initial evidence be the set of QRS annotations obtained by a state-of-the art detection algorithm [32]. In this short strip, the eighth and ninth annotations correspond to false positives caused by the presence of noise. A classification-based strategy processes these two annotations as true QRS complexes, and the monotone nature of the reasoning prevents their possible refutation, probably leading to beat misclassification and false arrhythmia detection, propagating the errors onwards to the end of the processing. In contrast, the present framework provides a single normal rhythm as the best interpretation, which explains all but those two aforementioned annotations, being ignored and considered unintelligible from the available model. It is also worth noting the ability of this framework to integrate the results of an available classifier -as a sort of constraint specification- in the interpretation cycle.

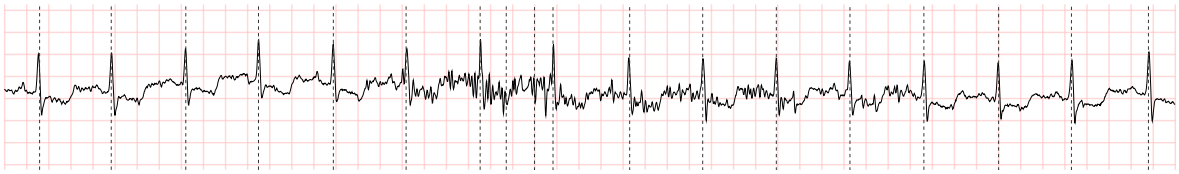

Figure 9: Unintelligible evidence due to noise. [Source: MIT-BIH arrhythmia DB, recording: 112, between 13:46.200 and 13:56.700]

## 6.3 Looking for missing evidence

Application of the classification paradigm to pattern detection also entails the potential risk of providing false negative results. In the worst case, a false negative result may be interpreted by a decision maker as an evidence of absence, leading to interpretation errors with their subsequent costs; or in the best case, as an absence of evidence caused by the lack of a proper detection instrument.

Even though abduction is fallible, and false negative results persist, the hypothesize-and-test cycle involves a prediction mechanism that points to that missing evidence that should be expected, and moreover, estimates when it should appear. Both the bottom-up and the top-down processing performed along this cycle reinforces the confidence in the interpretation, since the semantics of any conclusion is widened according to its explanatory power.

As an example, consider the interpretation problem illustrated in figure 10. The initial evidence is again a set of QRS annotations obtained by a state-of-the art detection algorithm [32]. Note that the eighth beat has not been annotated, due to a sudden decrease in the signal amplitude. This error can be amended along the hypothesis-and-test cycle, since the normal rhythm hypothesis that abstracts the first seven QRS annotations predicts the following QRS to be in the position of the missing annotation, and the PREDICT() procedure can look for it (e.g. checking an alternative set of constraints).

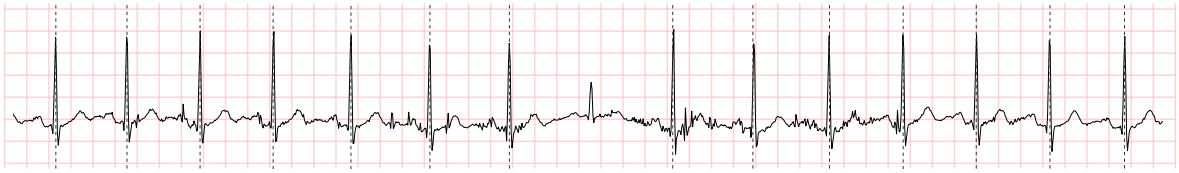

Figure 10: Missing evidence that may be discovered by prediction. [Source: MIT-BIH normal sinus rhythm DB, recording: 18184, between 09:12:45.000 and 09:12:55.500]

The capability of abduction to ignore or look for new evidence has been tested with a simplified version of the present framework in the QRS detection problem [30], leading to a statistically significant improvement over a state-of-the art algorithm.

## 6.4 Interpretability of the reasoning process and the results

Interpretability of a reasoning formalism, defined as the ability to understand and evaluate its conclusions, is an essential feature to achieve an adequate confidence for decision making [22]. In this sense, there is a number of classification methods with good interpretability, but methods that typically offer the best performance belong to the so-called *black box* approaches.

The present interpretation framework is able to provide a justification of any result in relation to the available model. Given any solution or partial solution of an interpretation problem, the searching path up to $I_0$ gives us a full detail of all those reasoning steps taken to this end, and any abstraction hypothesis can be traced back to the information supporting it.

This interpretation framework is also able to answer why a certain hypothesis has been rejected or neglected at any reasoning step. This is done by exploring those branches

outside the path between $I_0$ and the solution. Since the $K$ exploration parameter within the CONSTRUE() algorithm has been chosen as the maximum number of hypotheses that may explain a given observable, it is possible to reproduce the reasoning steps to conjecture any abstraction hypothesis, and check why it did not succeed (non-satisfaction of pattern constraints, lower coverage, etc.). This can be useful to build and refine the knowledge base.

# 7 Discussion

A new model-based framework for time series interpretation is proposed. This framework relies on some basic assumptions: (i) interpreting the behavior of a system from the set of available observations is a sort of conjecturing, and as such follows the logic of abduction; (ii) the interpretation task involves both bottom-up and top-down processing of information along a set of abstraction levels; (iii) at the lower levels of abstraction, the interpretation task is a form of precompiled knowledge-based pattern recognition; (iv) the interpretation task involves both the representation of time and the reasoning about time and along time.

Model-based representation in the present framework is based on the notion of abstraction pattern, which defines an abstraction relation between observables, and provides the knowledge and methods to conjecture new observations from previous ones. Let us deepen in both the backward and forward logical meaning of an abstraction pattern, following a reasoning similar to that of [2]:

- *Backward meaning*. From a backward reading of an abstraction pattern, a hypothesis $h$ is a possible abstraction of $m_1, \ldots, m_n$, provided that the constraints in $C$ hold. An abstraction pattern satisfies the *compositionality principle* of abductive reasoning, and hence an abstraction hypothesis can be conjectured from a single piece of evidence, and new pieces of evidence can be later added [12]. On the other hand, the abductive observation of $h$ must satisfy the *falsification principle*, that is, if any of the constraints in $C$ is inconsistent with the evidence, then $h$ is refuted assuming this pattern is the only knowledge available about observing $h$. If there are multiple ways of observing $h$ by means of multiple patterns, and all of their constraints are inconsistent with evidence, we do not conclude $\neg h$, interpreted as failure to prove $h$; we will conclude $\neg h$ in all those interpretations conjecturing an observation of a different $h'$, being $h$ and $h'$ mutually exclusive.

- *Forward meaning*. An abductive observation is built upon an archetypical representation of a hypothesis $h$, creating an observation as an instance of $h$ by estimating, from the available evidence, its attribute values $\mathbf{A}_h$ and its temporal location $T_h$ by means of $\pi()$. From a forward reading, assuming $h$ is true there is an observation for each observable of the set $m_1$,...,$m_n$ such that the constraints in $C$ hold. However, the estimated nature of abstraction does not usually allow us to infer, from the observation of $h$, the same observations of $m_1$,...,$m_n$ that have been abstracted into $h$. We must presume instead that assuming $h$ is true entails the occurrence of an observation for each observable of $m_1, ..., m_n$, without necessarily entailing its attribute values and its temporal location.

An abstraction model, built on a set of abstraction patterns, establishes a causal responsibility for the behavior observed in a complex system [17]. This responsibility is expressed in the language of processes: a process is said to be observable if it is assumed that it causes a recognizable trace in the physical quantity to be interpreted. The successive application of abstraction hypotheses provides successive reinterpretations of the observations, as more processes are included, expanding its scale and scope as we go up the abstraction hierarchy. This notion of causality is behind perception, concerned at the explanation of sensory data, in contrast with the notion of causality in diagnosis, concerned at the explanation of abnormality [6].

Representing and reasoning about context is a relevant issue in model-based diagnosis [2, 6, 24, 29]. A contextual observation is nothing more than another observation that needs not be explained by a diagnosis. In most of the works, the distinction between this two roles must be defined beforehand. Some other works enable the same observation to play different roles in different causal patterns, providing some general operations for expressing common changes made by the context in a diagnostic pattern [18, 23]. In the present interpretation framework an observation can either be part of the evidence to be explained in a certain abstraction pattern, or can be part of the environment in another abstraction pattern. Both sort of observations take part of the hypothesize-and-test cycle, with the only difference that observations from the environment of an abstraction pattern are not expected to be abstracted by this pattern. Hence, observations from the environment are naturally included in the deduction, subsumption and prediction modes of reasoning.

There are some issues where our efforts are focusing on at present. An important one is the representation of negated evidence and its processing by the CONSTRUE() algorithm. There can be certain observations that are incompatible with a certain interpretation, so their *absence* is a necessary condition. This can help us to model an *inhibitory relation* between processes, that allows us to represent that a certain process prevents other to occur under some temporal constraints.

Further efforts are being made to improve the efficiency of the interpretation process. To this end, two main strategies are being explored. On the one hand, the exploitation of the model structure to identify necessary and sufficient conditions for every hypotheses to be conjectured; necessary conditions avoid the expansion of those hypotheses that can be ruled out because they are inconsistent with observations, while sufficient conditions avoid the construction of redundant interpretations [5]. Another different strategy entails additional restrictions in the amount of computer memory and time needed to run the algorithm, resulting in a selective pruning of the node expansion while sacrificing optimality; this strategy is similar to the one used in the K-Beam algorithm [9].

The CONSTRUE() algorithm is based on the assumption that all the evidence to be explained is available at the beginning of the interpretation task. A new version of the algorithm should be provided to cope with a wide range of problems where the interpretation must be updated as new evidence becomes available over time. Examples of such problems are continuous biosignal monitoring or plan execution monitoring [1]. At the emergence of a new piece of evidence, two reasoning modes may come into play triggered by the CONSTRUE() algorithm: a new explanatory hypothesis can be conjectured by means of the ABDUCE() procedure, or the evidence can be enclosed in an existing hypothesis by means of the SUBSUME() procedure. In this way, the incorporation of new evidence over time is seamlessly integrated into the hypothesize-and-test cycle. Furthermore, to properly address

these interpretation scenarios, the heuristics used to guide the search must be updated to account for the timing of the interpretation process, which will lead to the definition of a covering ratio until time t, and the complexity until time t.

On the other hand, the present framework has to be tested in different interpretation problems. An example of a problem that fits this approach is *chronicle recognition* [1]. In this case, the initial evidence is a potentially large collection of time-stamped data, and the objective is to summarize it in a set of chronicles, describing more abstract events and relations. Notions like the focus of attention or the use of abstraction levels have already been used by other authors [8], albeit from a different, non abductive perspective. Currently, in the ECG domain, we are working on some significant open problems, like beat labeling, with some promising results which show a better performance than the previous state-of-the-art automatic machine learning classifiers, and a performance comparable to the assisted machine learning classifiers that benefit from the aid of an expert cardiologist.

# Implementation

With the aim of supporting reproducible research, the full source code of the algorithms presented in this paper has been published under an open source license[1], along with a knowledge base for the interpretation of the ECG signal strips of all examples in this paper.

# Acknowledgements

This work was supported by the Spanish Ministry of Economy and Competitiveness under project TIN2014-55183-R. T. Teijeiro was funded by an FPU grant from the Spanish Ministry of Education (MEC) (ref. AP2010-1012).

# References

[1] R. Barták, R. A. Morris, and K. B. Venable. An Introduction to Constraint-Based Temporal Reasoning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(1):1–121, feb 2014.

[2] V. Brusoni, L. Console, P. Terenziani, and D. Theseider Dupré. A spectrum of definitions for temporal model-based diagnosis. *Artificial Intelligence*, 102(1):39–79, 1998.

[3] E. Charniak. Motivation analysis, abductive unification and nonmonotonic equality. *Artificial Intelligence*, 34(3):275–295, 1989.

[4] C.K. Chow. On optimum recognition error and reject tradeoff. *IEEE Transaction on Information Theory*, 16(1):41–46, 1970.

[1] `https://github.com/citiususc/construe`

[5] L. Console, L. Portinale, and D. Theseider Dupré. Using compiled knowledge to guide and focus abductive diagnosis. *IEEE Transactions on Knowledge and Data Engineering*, 8(5):690–706, 1996.

[6] L. Console and P. Torasso. A spectrum of logical definitions of model-based diagnosis. *Computational Intelligence*, 3(7):133–141, 1991.

[7] Working Party CSE. Recommendations for measurement standards in quantitative electrocardiography. *European Heart Journal*, 6(10):815–825, 1985.

[8] C. Dousson and P. Le Maigat. Chronicle recognition improvement using temporal focusing and hierarchization. In *The International Joint Conference on Artificial Intelligence*, pages 324–329, 2007.

[9] S. Edelkamp and S. Schrödl. *Heuristic Search: Theory and Applications*. Morgan Kaufmann, 2011.

[10] A Felner, S Kraus, and RE Korf. KBFS: K-best-first search. *Annals of Mathematics and Artificial Intelligence*, 39(1-2):19–39, 2003.

[11] D. Ferrucci, A. Levas, S. Bagchi, D. Gondek, and E.T. Mueller. Watson: Beyond Jeopardy. *Artificial Intelligence*, 199–200:93–105, 2012.

[12] P. Flach. Abduction and induction: Syllogistic and inferential perspectives. In *Abductive and Inductive Reasoning Workshop Notes*, pages 31–35. University of Bristol, 1996.

[13] G. Fumera, F. Roli, and G. Giacinto. Reject option with multiple thresholds. *Pattern Recognition*, (33):2099–2101, 2000.

[14] A. L. Goldberger et al. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation*, 101(23):215–220, June 2000.

[15] C. Hartshorn et al. *Collected papers of Charles Sanders Peirce*. Harvard University Press, 1931.

[16] J.R. Hobbs, M. Stickel, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63:69–142, 1993.

[17] J.R. Josephson and S.G. Josephson. *Abductive inference. Computation, philosophy, technology*. Cambridge University Press, 1994.

[18] J. M. Juarez, M. Campos, J. Palma, and R. Marin. Computing context-dependent temporal diagnosis in complex domains. *Expert Systems with Applications*, 35(3):991–1010, oct 2008.

[19] P. Laguna, R. Jané, and P. Caminal. Automatic detection of wave boundaries in multilead ECG signals: validation with the CSE database. *Computers and Biomedical Research*, 27:45–60, 1994.

[20] D. Litman and J. Allen. A plan recognition model for subdialogues in conversation. *Cognitive Science*, 11:163–200, 1987.

[21] F. Mörchen. Time series feature extraction for data mining using DWT and DFT. Technical Report no. 33, Department of Mathematics and Computer Science, University of Marburg, 2003.

[22] D. Nauck and R. Kruse. Obtaining interpretable fuzzy classification rules from medical data. *Artificial Intelligence in Medicine*, 16(2):149–169, jun 1999.

[23] J. Palma, J. M. Juarez, M. Campos, and R. Marin. Fuzzy theory approach for temporal model-based diagnosis: An application to medical domains. *Artificial Intelligence in Medicine*, 38(2):197, 2006.

[24] Y. Peng and J.A. Reggia. *Abductive inference models for diagnostic problem-solving*. Springer-Verlag, 1990.

[25] A. Petrenas, V. Marozas, and L. Sörnmo. Low-complexity detection of atrial fibrillation in continuous long-term monitoring. *Computers in biology and medicine*, 65:184–91, oct 2015.

[26] M.A.F. Pimentel, D.A. Clifton, L. Clifton, and L. Tarassenko. A review of novelty detection. *Signal Processing*, (99):215–249, 2014.

[27] D. Poole. A methodology for using a default and abductive reasoning system. *International Journal of Intelligent Systems*, 5(5):521–548, 1990.

[28] L. Sacchi, E. Parimbelli, S. Panzarasa, N. Viani, E. Rizzo, C. Napolitano, R. Ioana Budasu, and S. Quaglini. Combining Decision Support System-Generated Recommendations with Interactive Guideline Visualization for Better Informed Decisions. In *Artificial Intelligence in Medicine*, pages 337–341. Springer International Publishing, 2015.

[29] Y Shahar. A framework for knowledge-based temporal abstraction. *Artificial intelligence*, 90(1–2):79–133, 1997.

[30] T. Teijeiro, P. Felix, and J. Presedo. Using Temporal Abduction for Biosignal Interpretation: A Case Study on QRS Detection. In *2014 IEEE International Conference on Healthcare Informatics*, pages 334–339. IEEE, September 2014.

[31] Galen S. Wagner. *Marriott's Practical Electrocardiography*. Lippincott Williams & Wilkins, 11 edition, 2008.

[32] W. Zong, G.B. Moody, and D. Jiang. A robust open-source algorithm to detect onset and duration of QRS complexes. In *Computers in Cardiology*, pages 737–740, 2003.